

Encyclopedia of
Complexity and Systems Science Series
Editor-in-Chief: Robert A. Meyers

SPRINGER
REFERENCE

Marilda Sotomayor · David Pérez-Castrillo
Filippo Castiglione *Editors*

Complex Social and Behavioral Systems

Game Theory and Agent-Based Models

A Volume in the Encyclopedia of
Complexity and Systems Science,
Second Edition

 Springer

Encyclopedia of Complexity and Systems Science Series

Editor-in-Chief

Robert A. Meyers

The Encyclopedia of Complexity and Systems Science Series of topical volumes provides an authoritative source for understanding and applying the concepts of complexity theory together with the tools and measures for analyzing complex systems in all fields of science and engineering. Many phenomena at all scales in science and engineering have the characteristics of complex systems and can be fully understood only through the transdisciplinary perspectives, theories, and tools of self-organization, synergetics, dynamical systems, turbulence, catastrophes, instabilities, nonlinearity, stochastic processes, chaos, neural networks, cellular automata, adaptive systems, genetic algorithms, and so on. Examples of near-term problems and major unknowns that can be approached through complexity and systems science include: the structure, history, and future of the universe; the biological basis of consciousness; the integration of genomics, proteomics, and bioinformatics as systems biology; human longevity limits; the limits of computing; sustainability of human societies and life on earth; predictability, dynamics, and extent of earthquakes, hurricanes, tsunamis, and other natural disasters; the dynamics of turbulent flows; lasers or fluids in physics; microprocessor design; macromolecular assembly in chemistry and biophysics; brain functions in cognitive neuroscience; climate change; ecosystem management; traffic management; and business cycles. All these seemingly diverse kinds of phenomena and structure formation have a number of important features and underlying structures in common. These deep structural similarities can be exploited to transfer analytical methods and understanding from one field to another. This unique work will extend the influence of complexity and system science to a much wider audience than has been possible to date.

More information about this series at <https://link.springer.com/bookseries/15581>

Marilda Sotomayor
David Pérez-Castrillo • Filippo Castiglione
Editors

Complex Social and Behavioral Systems

Game Theory and Agent-Based
Models

A Volume in the Encyclopedia of Complexity
and Systems Science, Second Edition

With 158 Figures and 17 Tables

 Springer

Editors

Marilda Sotomayor
Department of Economics
University of Sao Paulo
Sao Paulo, Brazil

EPGE Brazilian School of Economics
and Finance
Sao Paulo, Brazil

David Pérez-Castrillo
Dept. of Economics and Economic
History
Universitat Autònoma de Barcelona and
Barcelona GSE
Barcelona, Spain

Filippo Castiglione
Istituto Applicazioni del Calcolo (IAC)
Consiglio Nazionale delle Ricerche (CNR)
Rome, Italy

ISBN 978-1-0716-0367-3 ISBN 978-1-0716-0368-0 (eBook)
ISBN 978-1-0716-0369-7 (print and electronic bundle)
<https://doi.org/10.1007/978-1-0716-0368-0>

© Springer Science+Business Media, LLC, part of Springer Nature 2020

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors, and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Science+Business Media, LLC, part of Springer Nature.

The registered company address is: 1 New York Plaza, New York, NY 10004, U.S.A.

Series Preface

The Encyclopedia of Complexity and System Science Series is a multivolume authoritative source for understanding and applying the basic tenets of complexity and systems theory as well as the tools and measures for analyzing complex systems in science, engineering, and many areas of social, financial, and business interactions. It is written for an audience of advanced university undergraduate and graduate students, professors, and professionals in a wide range of fields who must manage complexity on scales ranging from the atomic and molecular to the societal and global.

Complex systems are systems that comprise many interacting parts with the ability to generate a new quality of collective behavior through selforganization, e.g., the spontaneous formation of temporal, spatial, or functional structures. They are therefore adaptive as they evolve and may contain self-driving feedback loops. Thus, complex systems are much more than a sum of their parts. Complex systems are often characterized as having extreme sensitivity to initial conditions as well as emergent behavior that are not readily predictable or even completely deterministic. The conclusion is that a reductionist (bottom-up) approach is often an incomplete description of a phenomenon.

This recognition that the collective behavior of the whole system cannot be simply inferred from the understanding of the behavior of the individual components has led to many new concepts and sophisticated mathematical and modeling tools for application to many scientific, engineering, and societal issues that can be adequately described only in terms of complexity and complex systems.

Examples of Grand Scientific Challenges which can be approached through complexity and systems science include: the structure, history, and future of the universe; the biological basis of consciousness; the true complexity of the genetic makeup and molecular functioning of humans (genetics and epigenetics) and other life forms; human longevity limits; unification of the laws of physics; the dynamics and extent of climate change and the effects of climate change; extending the boundaries of and understanding the theoretical limits of computing; sustainability of life on the earth; workings of the interior of the earth; predictability, dynamics, and extent of earthquakes, tsunamis, and other natural disasters; dynamics of turbulent flows and the motion of granular materials; the structure of atoms as expressed in the Standard Model and the formulation of the Standard Model and gravity into a Unified Theory; the

structure of water; control of global infectious diseases; and also evolution and quantification of (ultimately) human cooperative behavior in politics, economics, business systems, and social interactions. In fact, most of these issues have identified nonlinearities and are beginning to be addressed with nonlinear techniques, e.g., human longevity limits, the Standard Model, climate change, earthquake prediction, workings of the earth's interior, natural disaster prediction, etc.

The individual complex systems mathematical and modeling tools and scientific and engineering applications that comprised the *Encyclopedia of Complexity and Systems Science* are being completely updated and the majority will be published as individual books edited by experts in each field who are eminent university faculty members.

The topics are as follows:

Agent Based Modeling and Simulation
 Applications of Physics and Mathematics to Social Science
 Cellular Automata, Mathematical Basis of
 Chaos and Complexity in Astrophysics
 Climate Modeling, Global Warming, and Weather Prediction
 Complex Networks and Graph Theory
 Complexity and Nonlinearity in Autonomous Robotics
 Complexity in Computational Chemistry
 Complexity in Earthquakes, Tsunamis, and Volcanoes, and Forecasting and Early Warning of Their Hazards
 Computational and Theoretical Nanoscience
 Control and Dynamical Systems
 Data Mining and Knowledge Discovery
 Ecological Complexity
 Ergodic Theory
 Finance and Econometrics
 Fractals and Multifractals
 Game Theory
 Granular Computing
 Intelligent Systems
 Nonlinear Ordinary Differential Equations and Dynamical Systems
 Nonlinear Partial Differential Equations
 Percolation
 Perturbation Theory
 Probability and Statistics in Complex Systems
 Quantum Information Science
 Social Network Analysis
 Soft Computing
 Solitons
 Statistical and Nonlinear Physics
 Synergetics
 System Dynamics
 Systems Biology

Each entry in each of the Series books was selected and peer reviews organized by one of our university-based book Editors with advice and consultation provided by our eminent Board Members and the Editor-in-Chief.

This level of coordination assures that the reader can have a level of confidence in the relevance and accuracy of the information far exceeding than that generally found on the World Wide Web. Accessibility is also a priority and for this reason each entry includes a glossary of important terms and a concise definition of the subject. In addition, we are pleased that the mathematical portions of our Encyclopedia have been selected by Math Reviews for indexing in MathSciNet. Also, ACM, the world's largest educational and scientific computing society, recognized our *Computational Complexity: Theory, Techniques, and Applications* book, which contains content taken exclusively from the *Encyclopedia of Complexity and Systems Science*, with an award as one of the notable Computer Science publications. Clearly, we have achieved prominence at a level beyond our expectations, but consistent with the high quality of the content!

Palm Desert, CA, USA
July 2020

Robert A. Meyers
Editor-in-Chief

Volume Preface

Game theory is the study of decision problems which involve several individuals (the decision-makers or players) interacting rationally. The models of game theory are abstract representations of a number of real-life situations and have applications to economics, political sciences, computer sciences, evolutionary biology, social psychology, and law, among others. These applications are also important for the development of the theory, since the questions that emerge may lead to new theoretic results.

This volume provides the main features of Game Theory, covering most of the fundamental theoretical aspects under the cooperative, non-cooperative, and “general” or “mixed” approaches.

The cooperative approach focuses on the possible outcomes of the decision-makers’ interaction by abstracting from the actions or decisions that may lead to these outcomes. Specifically, cooperative game theory studies the interactions among coalitions of players. Its main question is: Given the sets of feasible payoffs for each coalition, what payoff will be awarded to each player? One can take a positive or normative approach to answering this question, and different solution concepts in the theory lead towards one or the other.

The non-cooperative approach focuses on the actions that the decision-makers can take. As argued by John von Neumann and Oskar Morgenstern in their famous 1944 book titled *Theory of Games and Economic Behavior*, most economic questions should be analyzed as games. Some games are dynamic, stressing the sequential nature of the various decisions that agents can make. Other situations are better modeled as static games.

The volume also considers contributions of game theory to mechanism design, which has helped the development of other key research areas such as auction theory, contract theory, and two-sided matching theory. Given the importance of these areas in game theory and in economics, several chapters are devoted to their study. The reader can also appreciate the many applications of game theory to practical problems in several contributions to this volume.

Finally, a section is dedicated to the modeling and simulation paradigm known as agent-based modeling (ABM) that is markedly useful in studying complex systems made up of a large number of interdependent objects. This paradigm is relatively immature, even though commonly applied in a broad spectrum of disciplines (game theory included), thus a clear-cut and widely accepted definition of high-level concepts of agents, environment, interactions, and so on is still lacking. This section addresses the epistemological

issues related to the agent-based paradigm of modeling of complex systems in order to attempt to reach a more general comprehension of emergent properties which, though ascribed to the definition of a specific application domain, are also universal.

The editors wish to thank the authors for their generosity in providing very valuable and timely contributions as well as their high motivation and strong engagement. We would also like to thank the publisher for the positive and straightforward collaboration, especially Meghna Singh and Neha Thapa for their support throughout the production of this volume.

Rio de Janeiro, Brazil
Barcelona, Spain
Rome, Italy
July 2020

Marilda Sotomayor
David Pérez-Castrillo
Filippo Castiglione

Contents

Part I Game Theory	1
Game Theory, Introduction to	3
Marilda Sotomayor and David Pérez-Castrillo	
Cooperative Games (Von Neumann-Morgenstern Stable Sets) ...	9
Ryo Kawasaki, Jun Wako and Shigeo Muto	
Cooperative Games	49
Roberto Serrano	
Dynamic Games with an Application to Climate Change Models	61
Prajit K. Dutta	
Static Games	83
Oscar Volij	
Correlated Equilibria and Communication in Games	107
Françoise Forges	
Bayesian Games: Games with Incomplete Information	119
Shmuel Zamir	
Repeated Games with Complete Information	139
Olivier Gossner and Tristan Tomala	
Repeated Games with Incomplete Information	157
Jérôme Renault	
Reputation Effects	185
George J. Mailath	
Zero-Sum Two Person Games	199
T. E. S. Raghavan	
Stochastic Games	229
Yehuda John Levy and Eilon Solan	
Signaling Games	251
Joel Sobel	

Inspection Games	269
Rudolf Avenhaus and Morton J. Canty	
Principal-Agent Models	287
Inés Macho-Stadler and David Pérez-Castrillo	
Differential Games	305
Marc Quincampoix	
Mechanism Design	317
Ron Lavi	
Auctions	335
Martin Pesendorfer	
Implementation Theory	349
Luis C. Corchón	
Two-Sided Matching Models	371
Marilda Sotomayor and Ömer Özak	
Market Design	401
Fuhito Kojima, Fanqi Shi and Akhil Vohra	
Cost Sharing in Production Economies	421
Maurice Koster	
Market Games and Clubs	463
Myrna Wooders	
Learning in Games	485
John Nachbar	
Fair Division	499
Steven J. Brams and Christian Klamler	
Social Choice Theory	511
Salvador Barberà	
Voting	529
Alvaro Sandroni, Jonathan Pogach, Michela Tincani, Antonio Penta and Deniz Selman	
Voting Procedures, Complexity of	543
Olivier Hudry	
Evolutionary Game Theory	573
William H. Sandholm	
Networks and Stability	609
Frank H. Page Jr. and Myrna Wooders	
Game Theory and Strategic Complexity	639
Kalyan Chatterjee and Hamid Sabourian	

Part II Agent-Based Models	659
Agent-Based Modeling and Simulation, Introduction to	661
Filippo Castiglione	
Agent-Based Modeling and Simulation	667
Stefania Bandini, Sara Manzoni and Giuseppe Vizzari	
Agent-Based Modeling, Mathematical Formalism for	683
Reinhard Laubenbacher, Abdul S. Jarrah, Henning S. Mortveit and S. S. Ravi	
Logic and Geometry of Agents in Agent-Based Modeling	705
Samson Abramsky	
Agent-Based Modeling and Artificial Life	725
Charles M. Macal	
Embodied and Situated Agents, Adaptive Behavior in	747
Stefano Nolfi	
Interaction-Based Computing in Physics	767
Franco Bagnoli	
Swarm Intelligence	791
Gerardo Beni	
Social Phenomena Simulation	819
Paul Davidsson and Harko Verhagen	
Agent-Based Computational Economics	825
Moshe Levy	
Cellular Automaton Modeling of Tumor Invasion	851
Haralambos Hatzikirou, Georg Breier and Andreas Deutsch	
Agent-Based Modeling and Computer Languages	865
Michael J. North and Charles M. Macal	
Computer Graphics and Games, Agent-Based Modeling in	891
Brian Mac Namee	
Agent-Based Modeling, Large-Scale Simulations	913
Hazel R. Parry	
Index	927

About the Editor-in-Chief



Robert A. Meyers

President: RAMTECH Limited

Manager, Chemical Process Technology, TRW Inc.

Post doctoral Fellow: California Institute of Technology

Ph.D. Chemistry, University of California at Los Angeles

B.A. Chemistry, California State University, San Diego

Biography

Dr. Meyers has worked with more than 20 Nobel laureates during his career and is the originator and serves as Editor-in-Chief of both the Springer Nature *Encyclopedia of Sustainability Science and Technology* and the related and supportive Springer Nature *Encyclopedia of Complexity and Systems Science*.

Education

Postdoctoral Fellow: California Institute of Technology

Ph.D. in Organic Chemistry, University of California at Los Angeles

B.A. Chemistry with minor in Mathematics, California State University, San Diego

Dr. Meyers holds more than 20 patents and is the author or Editor-in-Chief of 12 technical books including the *Handbook of Chemical Production Processes*, *Handbook of Synfuels Technology*, and *Handbook of Petroleum Refining Processes* now in 4th Edition, and the *Handbook of Petrochemical Production Processes*, now in its second edition, (McGraw-Hill) and the *Handbook of Energy Technology and Economics*, published by John Wiley & Sons; *Coal Structure*, published by Academic Press; and *Coal Desulfurization* as well as the *Coal Handbook* published by Marcel Dekker. He served as Chairman of the Advisory Board for *A Guide to Nuclear Power Technology*, published by John Wiley & Sons, which won the Association of American Publishers Award as the best book in technology and engineering.

About the Volume Editors



Marilda Sotomayor graduated in Mathematics in 1967 from the Federal University of Rio de Janeiro (UFRJ), received her master's degree in Mathematics in 1972 from the Institute of Pure and Applied Mathematics (IMPA), and completed her Ph.D. in Mathematics in 1981 from the Catholique University of Rio de Janeiro (PUC/RJ) and IMPA. She also obtained the Privat Dozen in 1999 from the University of São Paulo/SP (USP/SP). Before joining the University of São Paulo/SP in 1997, she taught at the Catholique University of Rio de Janeiro, University of the State of Rio de Janeiro, and Federal University of Rio de Janeiro. Dr. Sotomayor has worked as Distinguished Visiting Professor in the Department of Economics at Brown University (USA), Visiting Researcher at the University of Pittsburgh (USA), Institute des Hautes Etudes Scientifiques (France), Universitat Autònoma de Barcelona (Spain), Ecole Polytechnique (France), and University of California, Berkeley (USA). At present, she has retired from UFRJ and USP/SP and has a position at the Graduate School of Economics of the Getúlio Vargas Foundation-RJ. Her academic distinctions and prizes include the TWAS Prize, 2016, awarded by The World Academy of Sciences; her election for Fellow of the Econometric Society in 2003, for Fellow of the J. S. Guggenheim Foundation in 1993, for Fellow of the Game Theory Society in 2017, and for the Economic Theory Fellow in 2015 (awarded by the Society for the Advancement of Economic Theory); the Lanchester Prize 1990 (Mathematical of Operation Research), Haralambos Simeonides Prize 2001 (ANPEC), Mario Henrique Simonsen Prize 2006 (RBE/FGV), Adriano Romariz Duarte Prize 1996 (SBE); Medal

of Honour 2013 (Order of the Economists of Brazil – OEB); Medal of Honour, awarded by the Department of Economics of USP-SP, Brazil: “in recognition of the academic work of Marilda Sotomayor, from 1997 to 2014”; and Medal of Honour, awarded by the Graduate School of Economics of Getulio Vargas Foundation in honor of Marilda Sotomayor: “by her contribution to the Economic Sciences in the developing of the Matching Theory.” Dr. Sotomayor has contributed to the scientific community as Charter Member and Council Member of the Game Theory Society; Associate Editor of *Econometrica* and of the *Brazilian Review of Econometrics*; council member of the Econometric Society; guest Editor for the *International Journal of Game Theory* of the issue A collection of papers dedicated to David Gale on the occasion of his 85th birthday, published in 2008; Editor of the session of Game Theory of the Encyclopedia Complexity and Systems Sciences, published by Springer in 2009; and Associate Editor of the *International Journal of Game Theory* 2005–2018. She has also worked as Member of the Nominating Committee for Officers of the Econometric Society, Nominating Committee for Fellows of the Econometric Society, and as President of the Latin American Standing Committee of the Econometric Society.

Dr. Sotomayor’s work in the organization of conferences includes the scientific and local organization of three international workshops of the Game Theory Society, held at the University of São Paulo, in 2002, 2010, and 2014, respectively and the scientific organization of the 17th International Conference on Game Theory and Economic Applications and of the Gales Feast: A day in honor of the 85th birthday of David Gale, held at the State University of New York, Stony Brook, in 2006 and 2007, respectively.

She has also served as Member of the programme committee of several congresses (Latin American Meetings of the Econometric Society, World Congresses of the Econometric Society, the 3rd and 4th World Congress of the Game Theory Society).

Her research has been supported in part by grants from the CNPq. She has written the book *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis* with Alvin Roth, one of the winners of the Nobel Prize of Economics in 2012; three chapters of books; and over 50 peer-reviewed articles. The aforementioned book was awarded with the Lanchester Prize of 1990 and it was also honored with the celebration of the conference: Roth and Sotomayor: Twenty years after, held at Duke University, North Carolina, 2010, for celebrating the 20 years of publication of the book. Her area of interest is game theory, matching markets, and market design.



David Pérez-Castrillo earned a Ph.D. in economics from the Ecole des Hautes Etudes en Science Sociales, Paris. He had previously graduated in Mathematics from the University of the Basque Country in Bilbao. He is currently Professor at Universitat Autònoma de Barcelona and Research Professor at Barcelona GSE. He has been visiting professor in numerous universities, including Shanghai University of Finance and Economics, Waseda Institute for Advanced Study, Katholieke Universiteit Leuven, Paris School of Economics, University of Copenhagen, and University of California, San Diego.

He has been President of the Spanish Economic Association, President and General Secretary of the South-European Association of Economics, and Officer of the Executive Committee of the European Association for Research in Industrial Economics.

He is a member of the editorial board of the *Journal of Economics and Management Strategy*, *Economics Letters*, the *Journal of Public Economic Theory*, and the *International Game Theory Review*. He has also been managing editor of *Investigaciones Económicas*. His research on game theory and applied microeconomics has been published, among others, in *American Economic Review*, *Journal of Economic Theory*, *Journal of Financial Economics*, *Management Science*, *Games and Economic Behavior*, and *International Economic Review*. Because of his contributions, he has been awarded the Distinció per a la Promoció de la Recerca Universitària of the Generalitat de Catalunya for Young Researchers, ICREA Academia chairs, the Prize Haralambos Simeonidis, and the Arrow Price of the BE Journal.



Filippo Castiglione is Research Director at the Istituto per le Applicazioni del Calcolo of the National Research Council of Italy and adjunct professor of Machine Learning at the Department of Mathematics and Physics of Roma Tre University. He graduated in Computer Science at the University of Milan, Italy, and got a Ph.D. in Scientific Computing at the University of Cologne, Germany. Dr. Castiglione has been postdoc at the Institute for Medical Bio-Mathematics in Tel Aviv, Israel, and visiting research fellow at the IBM – T.J. Watson Research Center, Yorktown Heights (NY); in the Department of Molecular Biology at Princeton University; and in the Department of Cell Biology, Harvard Medical School, Boston. Dr. Castiglione has published one book and about 100 reviewed research papers among journals, books, and conferences proceedings. He is the main author of the C-ImmSim agent-based simulation model of the immune system. He has received funding from the EU in the FP6, FP7, H2020, IMI, and JPI Framework Programmes. His research interests range from the study of complex systems to the modeling of biological systems, machine learning, and high-performance computing.

Contributors

Samson Abramsky Department of Computer Science, University of Oxford, Oxford, UK

Rudolf Avenhaus Armed Forces University Munich, Neubiberg, Germany

Franco Bagnoli Department Physics and Astronomy and CSDC, University of Florence, Florence, Italy

Stefania Bandini Complex Systems and Artificial Intelligence Research Center, University of Milan-Bicocca, Milan, Italy

Salvador Barberà MOVE, Universitat Autònoma de Barcelona and Barcelona GSE, Barcelona, Spain

Gerardo Beni University of California, Riverside, CA, USA

Steven J. Brams Department of Politics, New York University, New York, NY, USA

Georg Breier Division of Medical Biology, Medical Faculty Carl Gustav Carus, Technische Universität Dresden, Dresden, Germany

Morton J. Canty Institute for Chemistry and Dynamics of the Geosphere, Forschungszentrum Jülich, Jülich, Germany

Filippo Castiglione Istituto Applicazioni del Calcolo (IAC), Consiglio Nazionale delle Ricerche (CNR), Rome, Italy

Kalyan Chatterjee Department of Economics, The Pennsylvania State University, University Park, USA

Luis C. Corchón Departamento de Economía, Universidad Carlos III, Madrid, Spain

Paul Davidsson Department of Computer Science, Malmö University, Malmö, Sweden

Andreas Deutsch Center for Information Services and High Performance Computing, Technische Universität Dresden, Dresden, Germany

Prajit K. Dutta Department of Economics, Columbia University, New York, NY, USA

Françoise Forges Ceremade, Université Paris-Dauphine, Paris, France

Olivier Gossner PSE, UMR CNRS-EHESS-ENPC-ENS 8545, Northwestern University, Paris, France

Haralambos Hatzikirou Center for Information Services and High Performance Computing, Technische Universität Dresden, Dresden, Germany
Helmholtz Centre for Infection Research, Department Systems Immunology, Braunschweig, Germany

Olivier Hudry École Nationale Supérieure des Télécommunications, Paris, France

Abdul S. Jarrah Department of Mathematics and Statistics, American University of Sharjah, Sharjah, United Arab Emirates

Ryo Kawasaki Department of Industrial Engineering and Economics, School of Engineering, Tokyo Institute of Technology, Tokyo, Japan

Christian Klamler Institute of Public Economics, University of Graz, Graz, Austria

Fuhito Kojima Department of Economics, Stanford University, Stanford, CA, USA

Maurice Koster University of Amsterdam, Amsterdam, The Netherlands

Reinhard Laubenbacher Virginia Bioinformatics Institute, Virginia Polytechnic Institute and State University, Virginia, USA

Ron Lavi The Technion – Israel Institute of Technology, Haifa, Israel

Moshe Levy The Hebrew University, Jerusalem, Israel

Yehuda John Levy Adam Smith Business School, University of Glasgow, Glasgow, UK

Charles M. Macal Center for Complex Adaptive Agent Systems Simulation (CAS2), Decision and Information Sciences Division, Argonne National Laboratory, Argonne, IL, USA

Inés Macho-Stadler Universitat Autònoma de Barcelona and Barcelona GSE, Barcelona, Spain

George J. Mailath Department of Economics, University of Pennsylvania, Philadelphia, USA

Sara Manzoni Complex Systems and Artificial Intelligence Research Center, University of Milan-Bicocca, Milan, Italy

Henning S. Mortveit Virginia Bioinformatics Institute, Virginia Polytechnic Institute and State University, Virginia, USA

Shigeo Muto Department of Business Economics, School of Management, Tokyo University of Science, Tokyo, Japan

John Nachbar Department of Economics, Washington University, St. Louis, MO, USA

Brian Mac Namee School of Computing, Dublin Institute of Technology, Dublin, Ireland

Stefano Nolfi Institute of Cognitive Sciences and Technologies, National Research Council (CNR), Rome, Italy

Michael J. North Argonne National Laboratory, Global Security Sciences Division, Argonne, IL, USA

Ömer Özak Department of Economics, Brown University, Providence, USA

Frank H. Page Jr. Department of Economics, Indiana University, Bloomington, IN, USA

Hazel R. Parry Central Science Laboratory, York, UK

Antonio Penta University of Pennsylvania, Philadelphia, PA, USA

David Pérez-Castrillo Dept. of Economics and Economic History, Universitat Autònoma de Barcelona and Barcelona GSE, Barcelona, Spain

Martin Pesendorfer Department of Economics, London School of Economics and Political Science, London, UK

Jonathan Pogach University of Pennsylvania, Philadelphia, PA, USA

Marc Quincampoix Laboratoire de Mathématiques de Bretagne Atlantique (LMBA), Université de Brest, Brest, France

T. E. S. Raghavan Department of Mathematics, Statistics and Computer Science, University of Illinois, Chicago, IL, USA

S. S. Ravi Department of Computer Science, University at Albany – State University of New York, New York, USA

Jérôme Renault Toulouse School of Economics, Université Toulouse 1 Capitole, Toulouse, France

Hamid Sabourian Faculty of Economics, University of Cambridge, Cambridge, UK

William H. Sandholm Department of Economics, University of Wisconsin, Madison, USA

Alvaro Sandroni University of Pennsylvania, Philadelphia, PA, USA

Deniz Selman University of Pennsylvania, Philadelphia, PA, USA

Roberto Serrano Department of Economics, Brown University, Providence, RI, USA

Fanqi Shi Department of Economics, Stanford University, Stanford, CA, USA

Joel Sobel Department of Economics, University of California, San Diego, CA, USA

Eilon Solan The School of Mathematical Sciences, Tel Aviv University, Tel Aviv, Israel

Marilda Sotomayor Department of Economics, University of Sao Paulo, Sao Paulo, Brazil

EPGE Brazilian School of Economics and Finance, Sao Paulo, Brazil

Michela Tincani University of Pennsylvania, Philadelphia, PA, USA

Tristan Tomala Economics and Finance Department, HEC Paris, Paris, France

Harko Verhagen Department of Computer and Systems Sciences, Stockholm University, Kista, Sweden

Giuseppe Vizzari Complex Systems and Artificial Intelligence Research Center, University of Milan-Bicocca, Milan, Italy

Akhil Vohra Department of Economics, Stanford University, Stanford, CA, USA

Oscar Volij Ben-Gurion University, Beer-Sheva, Israel

Jun Wako Department of Economics, Gakushuin University, Tokyo, Japan

Myrna Wooders Department of Economics, Vanderbilt University, Nashville, TN, USA

Shmuel Zamir Center for the Study of Rationality, Hebrew University, Jerusalem, Israel

Part I

Game Theory

Game Theory, Introduction to

Marilda Sotomayor^{1,2} and David Pérez-Castrillo³

¹Department of Economics, University of Sao Paulo, Sao Paulo, Brazil

²EPGE Brazilian School of Economics and Finance, Sao Paulo, Brazil

³Dept. of Economics and Economic History, Universitat Autònoma de Barcelona and Barcelona GSE, Barcelona, Spain

Game theory is the study of decision problems which involve several individuals (the decision-makers or players) interacting rationally. The models of game theory are abstract representations of a number of real-life situations and have applications to economics, political sciences, computer sciences, evolutionary biology, social psychology, and law, among others. These applications are also important for the development of the theory, since the questions that emerge may lead to new theoretic results.

This section is an attempt to provide the main features of game theory, covering most of the fundamental theoretical aspects under the cooperative, noncooperative, and “general” or “mixed” approaches.

The cooperative approach focuses on the possible outcomes of the decision-makers’ interaction by abstracting from the actions or decisions that may lead to these outcomes. Specifically, cooperative game theory studies the interactions among coalitions of players. Its main question is: Given the sets of feasible payoffs for each coalition, what payoff will be awarded to each player? One can take a positive or normative approach to answering this question and different solution concepts in the theory lead toward one or the other.

The first cooperative solution concept is the von Neumann-Morgenstern stable sets, treated in Chap. 1. However, the two best-known solution concepts in cooperative game theory are perhaps

the core and the Shapley value, which are presented and discussed in Chap. 2.

The noncooperative approach focuses on the actions that the decision-makers can take. Historically, the first contribution to the noncooperative game theory is due to Zermelo (1913), but the idea of a general theory of games was introduced by John von Neumann and Oskar Morgenstern in their famous book of 1944 entitled *Theory of Games and Economic Behavior*. These authors argued that most economic questions should be analyzed as games. They introduced the extensive-form and the strategic-form representations of a game, also known as dynamic and static games, respectively.

Dynamic games stress the sequentiality of the various decisions that agents can make. An essential component of a dynamic game is the description of who moves first, who moves second, etc. Static games, on the other hand, abstract from sequentiality of the possible moves and model interactions as simultaneous decisions. All extensive form games can be modeled as static games, and all strategic form games can be modeled as dynamic games. However some situations may be more conveniently modeled as one or the other kind of game. Dynamic games are examined in Chap. 3. The structure, as well as its principal results, is discussed in detail. The chapter ends with an important application, the economics of climate change.

The main ideas and results related to static games, as well as some interesting relationships that connect equilibrium concepts with the idea of rationality, are reviewed in Chap. 4. In this chapter, it is presented the general theorem of existence of strategic equilibria, due to Nash (1950). This result extends to more general games the minimax theorem, which was proved in von Neumann (1928) for two-player zero-sum games. In the literature there are two proofs published by Nash. One of them uses Brower’s fixed point theorem. The other one is a simpler proof, attributed to Gale by Nash, which uses Kakutani’s fixed point theorem. Some version of the proof that uses Brower’s fixed point theorem, by Geanakoplos

(2003), is presented in Chap. 4. Some discussion on correlated equilibrium and Bayesian games is also provided in this chapter.

The correlated equilibrium is a game theoretic solution concept proposed by Aumann (1974, 1987) in order to capture the strategic correlation opportunities that the players face when they take into account the extraneous environment in which they interact. Chapter 5 focuses on two possible extensions of the correlated equilibrium to Bayesian games: the strategic form correlated equilibrium and the communication equilibrium. The general framework of games with incomplete information is treated in Chap. 6, with special reference to “Bayesian games.”

Repeated games deal with situations in which a group of agents engage in a strategic interaction over and over. Chapter 7 is devoted to repeated games with complete information. In such games the data of the strategic interaction is fixed over time and is known by all the players. Chapter 8 discusses repeated games with incomplete information, a situation where several players repeat the same stage game, the players having different knowledge of the stage game which is repeated.

Repeated games have many equilibria, including the repetition of stage game Nash equilibria. At the same time, particularly when monitoring is imperfect, certain plausible outcomes are not consistent with equilibrium. Reputation effects is the term used for the impact upon the set of equilibria (typically of a repeated game) of perturbing the game by introducing incomplete information of a particular kind. This issue is treated in Chap. 9.

Games with two players are of particular significance. The first two-person game studied in the literature was the zero-sum two-person game, first analyzed by von Neumann and Morgenstern (1944). In such a game, one player’s gain is the other player’s loss. Chess, checkers, rummy, two finger morra, and tic-tac-toe are all examples of zero-sum two-person games. The theory for such games is surveyed in Chap. 10. Recent results on stochastic zero-sum games are presented in Chap. 11. Stochastic games are used to model dynamic interactions in which the environment changes in response to the behavior of the players. These games are discussed in Chap. 11.

Signaling games and inspection games are also two-player games. Signaling games is the subject of Chap. 12. They are games of incomplete information in which one player is informed and the other is not. Players can use the actions of their opponents to make inferences about the hidden information. The earliest work on this subject is Spence’s seminal 1972 work, in which education serves as a signal of ability. Inspection games are covered in Chap. 13. These games deal with the problem faced by an *inspector* who is required to control the compliance of an *inspectee* to some legal or otherwise formal undertaking. They started with the analysis of arms control and disarmament problems in the early 1960s and have been applied to auditing, environmental control, material accountancy, etc.

Inspections cause conflict in many real-world situations. In economics, there are services of many kinds, the fulfillment or payment of which has to be verified. One example is the problem of principal-agent relationships discussed in detail in Chap. 14. The principal-agent models provide the theory of contracts under asymmetric information, concerning relationships between owner and manager, insurer and insured, etc. The principal, e.g., an employer, delegates work or responsibility to the agent, the employee, and chooses a payment schedule that best exploits the agent’s self-interests. The agent, of course, behaves so as to maximize her own utility given the fee schedule proposed by the principal. The problem faced by the principal is to devise incentives to motivate the agent to act in the principal’s interest. This generates some type of transaction cost for the principal, which includes the task of investigating and selecting appropriate agents, gaining information to set performances standards, monitoring agents, bonding payments by the agents, and residual losses.

Chapter 15 is devoted to differential games with focus on two-player zero-sum and antagonist differential games. These are games in which the state of the players depends on time in a continuous way. The positions of the players are solutions to differential equations. Motivated by military applications in the “Cold War”, these games have a wide range of applications from economics to engineer sciences and recently to biology and behavioral ecology.

Mechanism designed is the subject of Chap. 16. It studies the construction of mechanisms that aim to reach a socially desirable outcome in the presence of rational but selfish players, who care only about their own private utility. More specifically, the question is how to design a mechanism such that the equilibrium behavior of the players in the game induced by the mechanism leads to the socially desired goal.

The theory of mechanism design has contributed to the development of other research areas as, for example, auction theory, contract theory, and two-sided matching theory. “For having laid the foundations of mechanism design theory” the 2007 Nobel Prize in Economics was awarded to Leonid Hurwicz, Eric Maskin, and Roger Myerson. Chapter 17 is devoted to the presentation of auctions and to introducing major contributions. It studies various auction formats, including English (ascending-price) and Dutch (descending-price) auctions, first-price and second-price sealed-bid auctions, as well as all-pay auctions.

A related theory is the theory of implementation, the subject of Chap. 18. It reverses the usual procedure, namely, fix a mechanism and see what the outcomes are. More precisely, it investigates the correspondence between normative goals and mechanisms designed to achieve those goals.

A class of “mixed” games is that of two-sided matching games, which has been analyzed since Gale and Shapley, 1962, under both cooperative and noncooperative game theoretic approaches. The two-sided matching theory is surveyed in Chap. 19 and 20. Chapter 19 focuses on the differences and similarities between some matching models. In their paper, Gale and Shapley formulated and solved the stable matching problem for the marriage and the college admissions markets. The solution of the college admissions problem was given by a simple deferred-acceptance algorithm which has been adapted and applied in the reorganization of admission processes of many two-sided matching markets. Chapter 20 studies the one-sided matching model and discusses applications, such as the medical residency matching, kidney exchange, and school choice.

Another class of problems that have been discussed from the perspective of cooperative and noncooperative game theory is the cost sharing problems, treated in Chap. 21. Applications are numerous ranging from environmental issues like pollution, fishing grounds, to sharing multipurpose reservoirs, road systems, communication networks, and the Internet. The worth of a “coalition” of such activities is defined as the hypothetical cost of carrying out the activities in that coalition only.

Market games and clubs are treated in Chap. 22 with focus on the equivalence between markets – defined as private goods economies where all participants in the economy have utility functions that are linear in the variable money – and games in characteristic function form.

Learning in games is surveyed in Chap. 23. It covers models in which players are “rational” but not necessarily in equilibrium: players forecast, possibly inaccurately, the future behavior of their opponents and optimize or ε optimize with respect to their forecasts.

Fair division is reviewed in Chap. 24. It provides a rigorous analysis of procedures for allocating goods, or deciding who wins on what issues, in a dispute.

Voting methods as ways to take collective decisions have been studied since ancient times. The contributions by Arrow (1951, 1963) and Black (1948, 1958) broaden the view by considering the design of collective-decision methods in general, from an axiomatic point of view. These procedures allow to aggregate preferences taking into account ethical and pragmatic principles, as well as the participants’ incentives. They are studied in Chap. 25.

The following two chapters deal with applications to political sciences. The first one, Chap. 26, presents a game theoretic analysis of voting systems as procedures to choose a winner among a set of candidates from the individual preferences of the voters or more ambitiously, allowing to rank all the candidates or a part of them. Such a situation occurs not only in the field of elections but also in many other fields as games, sports, artificial intelligence, spam detection, web search engines, and statistics. From a practical point of view, it is crucial to be able to announce who is the winner in a “reasonable” time. This raises the question of the

complexity of the voting procedures. The second chapter, Chap. 27, details the complexity results about several voting procedures.

Chapter 28 deals with applications to biology. This field, known as evolutionary game theory, started in 1972 with the publication of a series of papers by the mathematical biologist John Maynard Smith. Maynard Smith adapted the methods of traditional game theory, which were created to model the behavior of rational economic agents, to the context of biological natural selection.

Network models have a long history in sociology, natural sciences, and engineering. However, only recently economists have begun to think of political and economic interactions as network phenomena and to model everything as games of network formation. Chapter 29 is devoted to stable networks and the game theoretic underpinnings of stable networks.

Chapter 30 deals with some aspect of bounded rationality that has generated important work, namely, the presence of constraints on the capacities of players. Various constraints could be considered, for example, limits on the ability to plan ahead in intertemporal decision-making or on the ability to compute best responses. This chapter discusses cognitive costs to players of using strategies that depend on long histories of past play. This is done mainly in the context of bargaining and markets. It is shown that such complexity considerations often enable us to make sharp predictions. It is also considered the issue in the context of repeated games.

List of the Chapters

1. Cooperative Games (von Neumann-Morgenstern Stable Sets)
Authors: Ryo Kawasaki, Jun Wako and Shigeo Muto
2. Cooperative Games
Author: Roberto Serrano
3. Dynamic Games with an Application to Climate Change Models
Authors: Prajit K. Dutta
4. Static Games
Author: Oscar Volij
5. Correlated Equilibria and Communication in Games
Author: Françoise Forges
6. Bayesian Games: Games with Incomplete Information
Author: Shmuel Zamir
7. Repeated Games with Complete Information
Authors: Olivier Gossner and Tristan Tomala
8. Repeated Games with Incomplete Information
Author: Jérôme Renault
9. Reputation Effects
Authors: George Mailath
10. Zero-Sum Two Person Games
Author: T.E.S. Raghavan
11. Stochastic Games
Authors: Yehuda John Levy and Eilon Solan
12. Signaling Games
Author: Joel Sobel
13. Inspection Games
Authors: Rudolf Avenhaus and Morton J. Canty
14. Principal-Agent Models
Authors: David Pérez-Castrillo and Ines Macho-Stadler
15. Differential Games
Author: Marc Quincampoix
16. Mechanism Design
Author: Ron Lavi
17. Auctions
Author: Martin Pesendorfer
18. Implementation Theory
Author: Luis Corchon
19. Two-Sided Matching Models
Authors: Marilda Sotomayor and Ömer Özak
20. Market Design
Authors: Fuhito Kojima, Fanqi Shi and Akhil Vohra
21. Cost Sharing
Author: Maurice Koster
22. Market Games and Clubs
Author: Myrna Wooders
23. Learning in Games
Author: John Nachbar

-
- | | |
|---|---|
| <p>24. Fair Division
Authors: Steven Brams and Christian Klamler</p> <p>25. Social Choice Theory
Author: Salvador Barbera</p> <p>26. Voting
Authors: Alvaro Sandroni, Antonio Penta, Jonathan Pogach, Deniz Selman and Michela Tincani.</p> | <p>27. Voting Procedures, Complexity Of
Author: Olivier Hudry</p> <p>28. Evolutionary Game Theory
Author: William Sandholm</p> <p>29. Networks and Stability
Authors: Frank Page and Myrna Wooders</p> <p>30. Game Theory and Strategic Complexity
Authors: Kalyan Chatterjee and Hamid Sabourian</p> |
|---|---|

Cooperative Games (Von Neumann-Morgenstern Stable Sets)

Ryo Kawasaki¹, Jun Wako² and Shigeo Muto³

¹Department of Industrial Engineering and Economics, School of Engineering, Tokyo Institute of Technology, Tokyo, Japan

²Department of Economics, Gakushuin University, Tokyo, Japan

³Department of Business Economics, School of Management, Tokyo University of Science, Tokyo, Japan

Article Outline

Glossary

Definition of the Subject

Introduction

Stable Sets in Abstract Games

Stable Set and Core

Stable Sets in Characteristic Function Form Games

Applications of Stable Sets in Abstract and Characteristic Function Form Games

Farsighted Stable Sets in a General Setting

Future Directions

Bibliography

Glossary

Abstract game An abstract game consists of a set of outcomes and a binary relation, called domination, on the outcomes. Von Neumann and Morgenstern presented this game form for general applications of stable sets.

Characteristic function form game A characteristic function form game consists of a set of players and a characteristic function that gives each group of players, called a coalition, a value or a set of payoff vectors that they can gain by themselves. It is a typical representation of cooperative games. For characteristic

function form games, several solution concepts are defined such as von Neumann-Morgenstern stable set, core, bargaining set, kernel, nucleolus, and Shapley value.

Domination Domination is a binary relation defined on the set of imputations, outcomes, or strategy combinations, depending on the form of a given game. In characteristic function form games, an imputation is said to dominate another imputation if there is a coalition of players such that they can realize their payoffs in the former imputation by themselves and make each of them better off than in the latter. Domination given a priori in abstract games can be also interpreted in the same way. In strategic form games, domination is defined on the basis of commonly beneficial changes of strategies by coalitions.

External stability A set of imputations (outcomes, strategy combinations) satisfies external stability if any imputation outside the set is dominated by some imputation inside the set.

Farsighted stable set A farsighted stable set is a von Neumann-Morgenstern stable set defined by indirect domination. That is, a farsighted stable set satisfies internal stability and external stability with respect to indirect domination.

Imputation An imputation is a payoff vector in a characteristic function form game that satisfies group rationality and individual rationality. The former means that the players divide the amount that the grand coalition of all players can gain, and the latter says that each player is assigned at least the amount that he/she can gain by himself/herself.

Indirect domination Indirect domination is a domination relation that takes into account the farsightedness of the players. Intuitively, an outcome is indirectly dominated by another if the latter can be reached through a sequence of deviations, and the coalitions when making their moves in the deviations are made better off at the final outcome than at the outcome when they deviate.

Internal stability A set of imputations (outcomes, strategy combinations) satisfies internal stability if there is no domination between any two imputations in the set.

Strategic form game A strategic form game consists of a player set, each player's strategy set, and each player's payoff function. It is usually used to represent noncooperative games.

Von Neumann-Morgenstern stable set A set of imputations (outcomes, strategy combinations) is a von Neumann-Morgenstern stable set if it satisfies both internal and external stabilities.

Definition of the Subject

The von Neumann-Morgenstern stable set (hereafter stable set) is the first solution concept in cooperative game theory defined by J. von Neumann and O. Morgenstern. Though it was defined for cooperative games in characteristic function form, von Neumann and Morgenstern gave a more general definition of a stable set in abstract games. Later, J. Greenberg and M. Chwe cleared a way to apply the stable set concept to the analysis of noncooperative games in strategic and extensive forms. Stable sets in a characteristic function form game may not exist, as was shown by W. F. Lucas for a ten-person game that does not admit a stable set. On the other hand, stable sets exist in many important games. In voting games, for example, stable sets exist, and they indicate what coalitions can be formed in detail. The core, on the other hand, can be empty in voting games, though it is one of the best-known solution concepts in cooperative game theory. The analysis of stable sets is not necessarily straightforward, since it can reveal a variety of possibilities. However, stable sets give us deep insights into players' behavior, such as coalition formation, in economic, political, and social situations.

Introduction

For studies of economic or social situations where players can engage in cooperative behavior, the

stable set was defined by von Neumann and Morgenstern (1953) as a solution concept for characteristic function form cooperative games. They also defined the stable set in abstract games so that one can apply the concept to more general games including noncooperative situations. Greenberg (1990) and Chwe (1994) cleared a way to apply the stable set concept to the analysis of noncooperative games in strategic and extensive forms.

The stable set is a set of outcomes satisfying two stability conditions: internal and external stability. Internal stability states that between any two outcomes in the set, there is no group of players such that all of its members prefer one to the other and they can realize the preferred outcome. External stability states that for any outcome outside the set, there is a group of players such that all of its members have a commonly preferred outcome in the set and they can realize it. Though the existence of stable sets does not hold in general as was shown by Lucas (1968) and Lucas and Rabie (1982), the stable set has revealed many interesting behaviors of players in economic, political, and social systems.

Von Neumann and Morgenstern (and also Greenberg) took into account only a single move by a group of players and their ability to foresee subsequent moves made by other groups of players. Harsanyi (1974) first pointed out that stable sets in characteristic function form games may not take into account the farsighted behavior of players and their ability to foresee subsequent moves made by other groups of players. Harsanyi's work inspired Chwe (1994) to incorporate the notion of foresight in social environments into the von Neumann-Morgenstern stability.

Chwe focused on a possible chain of moves, where a move by a group of players will cause a sequence of moves from other groups of players. Then the group of players moving first should take into account the sequence of moves that may follow and evaluate their profits of the final outcome of the sequence rather than the outcomes in the intermediate steps of the sequence. By incorporating such a sequence of moves, Chwe (1994) defined a more forward-looking concept, which

we call a farsighted stable set in what follows. The incorporation of the farsighted stable set provides richer results than the myopic stable set in many classes of games.

The rest of the chapter is organized as follows. The first several sections cover the definitions of stable sets in abstract games and other more specific models in cooperative game theory. The examples include voting games, production games, assignment games, marriage games, and house barter games. The second set of sections then covers farsighted stable sets, which are defined in response to the criticism that the original stable set is myopic. This second part also starts with the definition of the concepts in an abstract setting and then proceeds to examples which include strategic form games, characteristic function form games, coalition formation games, matching games, and house barter games. The chapter then concludes with some remarks on the possible directions for future research.

Stable Sets in Abstract Games

An abstract game is a pair $(W, \succ)$ of a set of outcomes W and an irreflexive binary relation $\succ$ on W , where irreflexivity means that $x \succ x$ does not hold for any $x \in W$. The relation $\succ$ is interpreted as follows: if $x \succ y$ holds, then there must exist a set of players such that they can induce x from y by themselves and all of them are better off in x .

A subset K of W is called a stable set of abstract game $(W, \succ)$ if the following two conditions are satisfied:

1. Internal stability: For any two elements $x, y \in K$, $x \succ y$ never holds.
2. External stability: For any element $z \notin K$, there must exist $x \in K$ such that $x \succ z$.

We explain more in detail what the external and internal stability conditions imply in the definition of a stable set. Suppose that players have a common understanding that each outcome inside a stable set is “stable” and that each outcome outside the set is “unstable.” Here

“stability” means that no group of players has an incentive to deviate from it, and “instability” means that there is at least one group of players that has an incentive to deviate from it. Then the internal and external stability conditions guarantee that the common understanding is never disproved and thus continues to prevail. In fact, suppose that the set is both internally and externally stable and take any outcome in the set. Then by internal stability, no group of players can be better off by deviating from it and inducing another outcome inside the set. Thus, no group of players reaches an agreement to deviate, which makes each outcome inside the set remain stable. Deviating players may be better off by inducing an outcome outside the set, but outcomes outside the set are commonly considered unstable. Thus, deviating players can never expect that such an outcome will continue. Next take any outcome outside the set. Then by external stability, there exists at least one group of players who can become better off by deviating from it and inducing an outcome inside the set. The induced outcome is considered stable since it is in the set. Hence, the group of players will deviate to a stable outcome, thereby reinforcing that the outside outcomes are unstable.

Stable Set and Core

Another solution concept that is widely known is the core. For a given abstract game $G = (W, \succ)$, a subset C of W is called the core of G if $C = \{x \in W \mid \text{there is no } y \in W \text{ with } y \succ x\}$.

From the definition, the core satisfies internal stability. Thus, the core C of G is contained in any stable set of G if the latter exists. To see this, suppose that $C \not\subset K$ for a stable set K and C is nonempty, i.e., $C \neq \emptyset$. (If $C = \emptyset$, then clearly $C \subset K$.) Take any element $x \in C \setminus K$. Since $x \notin K$, by external stability, there exists $y \in K$ with $y \succ x$, which contradicts $x \in C$. When the core of a game satisfies external stability, it has very strong stability since the core itself is now a stable set. In this case, it is called the stable core. The stable core is the unique stable set of the game.

Stable Sets in Characteristic Function Form Games

An n -person game in characteristic function form with transferable utility is a pair (N, v) of a player set $N = \{1, 2, \dots, n\}$ and a characteristic function v on the set 2^N of all subsets of N such that $v(\emptyset) = 0$. Each subset of N is called a coalition. The game (N, v) is often called a transferable utility (TU) game. A characteristic function form game without transferable utility is called a non-transferable utility (NTU) game: its characteristic function gives each coalition a set of payoff vectors that the players in the coalition can get. For NTU games and their stable sets, refer to Aumann and Peleg (1960) and Peleg (1986). In this section, we will only deal with TU characteristic function form games and refer to them simply as characteristic function form games. Let (N, v) be a characteristic function form game. The characteristic function v assigns a real number $v(S)$ to each coalition $S \subseteq N$. The value $v(S)$ indicates the worth that coalition S can achieve by itself. An n -dimensional vector $x = (x_1, x_2, \dots, x_n)$ is called a payoff vector. A payoff vector x is called an imputation if the following two conditions are satisfied:

1. Group rationality: $\sum_{i=1}^n x_i = v(N)$.
2. Individual rationality: $x_i \geq v(\{i\})$ for each $i \in N$.

The first condition says that all players cooperate and share the worth $v(N)$ that they can produce. The second condition says that each player must receive at least the amount that he/she can gain by himself/herself. Let A be the set of all imputations. Let x, y be any two imputations and S be any coalition. We say that x dominates y via S and write this as $x \text{ dom}_S y$ if the following two conditions are satisfied:

1. Coalitional rationality: $x_i > y_i$ for each $i \in S$.
2. Effectivity: $\sum_{i \in S} x_i \leq v(S)$.

The first condition says that every member of coalition S strictly prefers x to y . The second condition says that coalition S can guarantee the payoff x_i for each member $i \in S$ by themselves.

We say that x dominates y (denoted by $x \text{ dom } y$) if there exists at least one coalition S such that $x \text{ dom}_S y$. It should be noted that a pair (A, dom) is an abstract game defined in the section “Stable Sets in Abstract Games.” It is easily seen that “dom” is an irreflexive binary relation on A . A stable set and the core of game (N, v) are defined as a stable set and the core of the associated abstract game (A, dom) , respectively. Since von Neumann and Morgenstern defined the stable set, whether one exists in a general game had been one of the most important problems in game theory. However, the problem was eventually solved in a negative way, as Lucas (1968) found the following ten-person characteristic function form game in which no stable set exists.

A game with no stable set: Consider the following ten-person game:

$$\begin{aligned}
 N &= \{1, 2, \dots, 10\}, \\
 v(N) &= 5, v(\{1, 3, 5, 7, 9\}) = 4, \\
 v(\{3, 5, 7, 9\}) &= v(\{1, 5, 7, 9\}) = v(\{1, 3, 7, 9\}) = 3, \\
 v(\{1, 2\}) &= v(\{3, 4\}) = v(\{5, 6\}) = v(\{7, 8\}) \\
 &= v(\{9, 10\}) = 1, \\
 v(\{3, 5, 7\}) &= v(\{1, 5, 7\}) = v(\{1, 3, 7\}) \\
 &= v(\{3, 5, 9\}) = v(\{1, 5, 9\}) = v(\{1, 3, 9\}) = 2, \\
 v(\{1, 4, 7, 9\}) &= v(\{3, 6, 7, 9\}) = v(\{5, 2, 7, 9\}) = 2, \\
 \text{and } v(S) &= 0 \text{ for all other } S \subset N.
 \end{aligned}$$

Though this game has no stable set, it has a nonempty core. A game with no stable set and an empty core was also found by Lucas and Rabie (1982). We remark on a class of games in which a stable core exists. As mentioned before, if a stable set exists, it always contains the core, which is of course true also in characteristic function form games. Furthermore, in characteristic function form games, there is an interesting class of games, called convex games, in which the core is a stable core. A characteristic function form game (N, v) is a convex game if for any $S, T \subseteq N$ with $S \subset T$ and for any $i \notin T$, $v(S \cup \{i\}) - v(S) \leq v(T \cup \{i\}) - v(T)$, i.e., the bigger coalition a player joins, the larger the player's contribution becomes. In convex games, the core is large and satisfies external stability. For the details, refer to Shapley (1971).

Applications of Stable Sets in Abstract and Characteristic Function Form Games

Symmetric Voting Games

This section deals with applications of stable sets to voting situations. Let us start with a simple example.

Example 1 Suppose there is a committee consisting of three players deciding on a bill. Each player has one vote and whether to pass the bill or not is decided by the simple majority rule. That is, to pass a bill, at least two votes are necessary. Before analyzing the players' behavior, we first formulate the situation as a characteristic function form game. Let the player set be $N = \{1, 2, 3\}$. Since a coalition of a simple majority of players can pass any bill, we give value 1 to such coalitions. Other coalitions can pass no bill. We thus give them value 0. Hence, the characteristic function is given by

$$v(S) = \begin{cases} 1 & \text{if } |S| \geq 2, \\ 0 & \text{if } |S| \leq 1, \end{cases}$$

where $|S|$ denotes the number of players in coalition S . The set of imputations is given by

$$A = \{x = (x_1, x_2, x_3) \mid x_1 + x_2 + x_3 = 1, x_1, x_2, x_3 \geq 0\}$$

One stable set of this game is given by the set K consisting of three imputations: $(1/2, 1/2, 0)$, $(1/2, 0, 1/2)$, and $(0, 1/2, 1/2)$. A brief proof is the following. Since each of the three imputations has only two numbers $1/2$ and 0 , internal stability is trivial. To show external stability, take any imputation $x = (x_1, x_2, x_3)$ from outside K . Suppose first $x_1 < 1/2$. Since $x \notin K$, at least one of x_2 and x_3 is less than $1/2$. If $x_2 < 1/2$, then $(1/2, 1/2, 0)$ dominates x via coalition $\{1, 2\}$. A similar argument can be used for when $x_3 < 1/2$ by using $(1/2, 0, 1/2)$ to dominate x via $\{1, 3\}$. Next suppose $x_1 = 1/2$. Since $x \notin K$, $0 < x_2, x_3 < 1/2$. Thus, $(0, 1/2, 1/2)$ dominates x via coalition $\{2, 3\}$. Finally suppose $x_1 > 1/2$. Then $x_2, x_3 < 1/2$, and thus, $(0, 1/2, 1/2)$ dominates x via coalition $\{2, 3\}$.

Thus, the proof of external stability is complete. This three-point stable set indicates that a two-person coalition is formed and that players in the coalition share equally the outcome obtained by passing a bill.

This game has three other types of stable sets. First, any set $K_c^1 = \{x \in A \mid x_1 = c\}$ with $0 \leq c < 1/2$ is a stable set. The internal stability of each K_c^1 is trivial. To show external stability, take any imputation $x = (x_1, x_2, x_3) \notin K_c^1$. Suppose $x_1 > c$. Define $y = (y_1, y_2, y_3)$ by $y_1 = c$, $y_2 = x_2 + (x_1 - c)/2$, $y_3 = x_3 + (x_1 - c)/2$. Then $y \in K_c^1$ and $y \text{ dom}_{\{2,3\}} x$. Next suppose $x_1 < c$. Notice that at least one of x_2 and x_3 is less than $1 - c$ since $c < 1/2$. Suppose without loss of generality $x_2 < 1 - c$. Since $c < 1/2$, we have $(c, 1 - c, 0) \in K_c^1$ and $(c, 1 - c, 0) \text{ dom}_{\{1,2\}} x$. Thus, external stability holds. This stable set indicates that player 1 gets a fixed amount c and players 2 and 3 negotiate for how to allocate the rest $1 - c$. Similarly, any sets $K_c^2 = \{x \in A \mid x_2 = c\}$ and $K_c^3 = \{x \in A \mid x_3 = c\}$ with $0 \leq c < 1/2$ are stable sets. The three-person game of Example 1 has no other stable set. See von Neumann and Morgenstern (1953). The first type of stable set (K) is called a symmetric (or objective) stable set, while the other types (K_c^1, K_c^2, K_c^3) are called discriminatory stable sets. As a generalization of the above result, symmetric stable sets are found in general n -person simple majority voting games. An n -person characteristic function form game (N, v) with $N = \{1, 2, \dots, n\}$ is called a simple majority voting game if

$$v(S) = \begin{cases} 1 & \text{if } |S| > n/2, \\ 0 & \text{if } |S| \leq n/2. \end{cases}$$

A coalition S with $v(S) = 1$, i.e., with $|S| > n/2$, is called a winning coalition. A winning coalition S is said to be minimal if $v(T) = 0$ for every strict subset T of S . In simple majority voting games, a minimal winning coalition is a coalition of $(n + 1)/2$ players if n is odd or $(n + 2)/2$ players if n is even. The following theorem holds. See Bott (1953) for the proof.

Theorem 1 Let (N, v) be a simple majority voting game. Then the following hold:

1. If n is odd, then the set

$$K = \underbrace{\langle 2/(n+1), \dots, 2/(n+1) \rangle}_{\frac{n+1}{2}}, \underbrace{\langle 0, \dots, 0 \rangle}_{\frac{n-1}{2}}$$

is a stable set where the symbol $\langle x \rangle$ denotes the set of all imputations obtained from x through permutations of its components.

2. If n is even, the set

$$K = \{ \{x \in A \mid \underbrace{x_1 \dots x_{n/2}}_{\frac{n}{2}} \geq \underbrace{x_{(n/2)+1} \dots x_n}_{\frac{n}{2}} \} \}$$

is a stable set, where

$$A = \left\{ x = (x_1, \dots, x_n) \mid \sum_{i=1}^n x_i = 1, x_1, \dots, x_n \geq 0 \right\}$$

and $\langle Y \rangle = \bigcup_{x \in Y} \langle x \rangle$ for a set Y .

It should be noted from the first proposition of Theorem 1 that when the number of players is odd, a minimal winning coalition is formed. The members of the coalition share equally the total profit. On the other hand, when the number of players is even, the second proposition of Theorem 1 shows that every player may gain a positive profit. This implies that the grand coalition of all players is formed. In negotiating on how to share the profit, two coalitions, each with $n/2$ players, are formed, and profits are shared equally within each coalition. Since at least $n/2 + 1$ players are necessary to win when n is even, an $n/2$ -player coalition is the smallest coalition that can prevent its complement from winning. Such a coalition is called a minimal blocking coalition. When n is odd, an $(n+1)/2$ -player minimal winning coalition is also a minimal blocking coalition.

General Voting Games

In this section, we present properties of stable sets and cores in general (not necessarily symmetric) voting games. A characteristic function form game (N, v) is called a simple game if $v(S) = 1$

or 0 for each nonempty coalition $S \subseteq N$. A coalition S with $v(S) = 1$ (resp. $v(S) = 0$) is a winning coalition (resp. losing coalition). A simple game is called a voting game if it satisfies the following: (1) $v(N) = 1$; (2) if $S \subseteq T$, then $v(S) \leq v(T)$; and (3) if S is winning, then $N \setminus S$ is losing. The first condition implies that the grand coalition N is always winning. The second condition says that a superset of a winning coalition is also winning. The third condition says that there are no two disjoint winning coalitions. It is easily shown that the simple majority voting game studied in the previous section satisfies these conditions. A player has veto power if he/she belongs to every winning coalition. As for cores of voting games, the following theorem holds.

Theorem 2 Let (N, v) be a voting game. Then the core of (N, v) is nonempty if and only if there exists a player who has veto power.

Thus, the core is not a useful tool for analyzing voting situations with no veto player. In simple majority voting games, no player has veto power, and thus, the core is empty. The following theorem shows that stable sets always exist.

Theorem 3 Let (N, v) be a voting game. Let S be a minimal winning coalition and define a set K by

$$K = \left\{ x \in A \mid \sum_{i \in S} x_i = 1, x_i = 0 \ \forall i \notin S \right\}.$$

Then K is a stable set.

Thus, in voting games, a minimal winning coalition is always formed, and they gain all the profit. For the proofs of these theorems, see Owen (1995). Further results on stable sets in voting games are found in Bott (1953), Griesmer (1959), Heijmans (1991), Lucas et al. (1982), Muto (1979, 1982b), Owen (1965), Rosenmüller (1977), and Shapley (1962, 1964).

Production Market Games

In this section, we consider the market production game and its stable sets, which are covered in

detail in Hart (1973) and Muto (1982a). To facilitate the discussion, let us start with a simple example.

Example 2 There are four players, each having one unit of a raw material. Two units of the raw material are necessary for producing one unit of an indivisible commodity. One unit of the commodity is sold at p dollars. The situation is formulated as the following characteristic function form game. The player set is $N = \{1, 2, 3, 4\}$. Since two units of the raw material are necessary to produce one unit of the commodity, the characteristic function v is given by

$$v(S) = \begin{cases} 2p & \text{if } |S| = 4, \\ p & \text{if } |S| = 3, 2, \\ 0 & \text{if } |S| = 1, 0. \end{cases}$$

The set of imputations is

$$A = \{x = (x_1, x_2, x_3, x_4) \mid x_1 + x_2 + x_3 + x_4 = 2p, \\ x_1, x_2, x_3, x_4 \geq 0\}.$$

The following set K is one of the stable sets of the game:

$$K = \langle \{x = (x_1, x_2, x_3, x_4) \in A \mid x_1 = x_2 = x_3 \geq x_4\} \rangle.$$

To show internal stability, take two imputations $x = (x_1, x_2, x_3, x_4)$ with $x_1 = x_2 = x_3 \geq x_4$ and $y = (y_1, y_2, y_3, y_4)$ in K . Suppose x dominates y . Since $x_1 = x_2 = x_3 \geq p/2 \geq x_4$, the domination must hold via coalition $\{i, 4\}$ with $i = 1, 2, 3$. Then we have a contradiction $2p = \sum_{i=1}^4 x_i > \sum_{i=1}^4 y_i = 2p$, since $y \in K$ implies that the largest three elements of y are equal. To show external stability, take $z = (z_1, z_2, z_3, z_4) \notin K$. Suppose $z_1 \geq z_2 \geq z_3 \geq z_4$. Then $z_1 > z_3$. Define $y = (y_1, y_2, y_3, y_4)$ by

$$y_i = \begin{cases} z_3 + \frac{z_1 + z_2 - 2z_3}{4} & \text{for } i = 1, 2, 3, \\ z_4 + \frac{z_1 + z_2 - 2z_3}{4} & \text{for } i = 4. \end{cases}$$

Then $y \in K$ and $y \text{ dom}_{\{3,4\}} z$, since $y_3 > z_3$, $y_4 > z_4$ and $y_3 + y_4 \leq p = v(\{3, 4\})$. This stable set

shows that in the negotiation of splitting the $2p$ dollars, three players form a coalition and share equally the gain obtained through collaboration. At least two players are necessary to produce the commodity. Thus, a three-player coalition is the smallest coalition that can prevent its complement from producing the commodity, i.e., a minimal blocking coalition. We would claim that in the market, a minimal blocking coalition is formed and that profits are shared equally within the coalition.

An extension of the model was given by Hart (1973) and Muto (1982a). Hart considered the following production market with n players, each holding one unit of a raw material. To produce one unit of an indivisible commodity, k units of raw materials are necessary. The associated production market game is defined by the player set $N = \{1, 2, \dots, n\}$ and the characteristic function v given by

$$v(S) = \begin{cases} 0 & \text{if } 0 \leq |S| < k, \\ p & \text{if } k \leq |S| < 2k, \\ \vdots & \\ jp & \text{if } jk \leq |S| < (j+1)k, \\ \vdots & \\ hp & \text{if } hk \leq |S| \leq n, \end{cases}$$

where $n = hk + r$ and h, r are integers such that $h \geq 1$ and $0 \leq r \leq k - 1$. When $h = 1$,

$$v(S) = \begin{cases} 0 & \text{if } |S| < k, \\ p & \text{if } |S| \geq k. \end{cases}$$

The following theorem holds.

Theorem 4 Suppose $h = 1$. Let $t = n - k + 1$ and $n = tu + w$ where u, w are integers such that $u \geq 1$ and $0 \leq w \leq t - 1$. Then the following set K is a stable set:

$$K = \langle \{x = (x_1, \dots, x_n) \in A \mid \\ x_1 = \dots = x_t \geq x_{t+1} = \dots = x_{2t} \\ \geq \dots \geq x_{tu+1} = \dots = x_n = 0\} \rangle,$$

where

$$A = \left\{ x = (x_1, \dots, x_n) \mid \sum_{i=1}^n x_i = p, x_1, \dots, x_n \geq 0 \right\}$$

is the set of imputations.

The theorem shows that in negotiating on how to share the profit, minimal blocking coalitions, i.e., coalitions of $n - k + 1$ players, are formed, and within each coalition, profits are shared equally. Players failing to form a coalition gain nothing. When $h \geq 2$, the following theorem holds.

Theorem 5 Suppose $h \geq 2$. Let

$$K = \left\langle \left\{ x = (x_1, \dots, x_n) \in A \mid x_1 = \dots = x_{n-k+1} \geq \frac{p}{k} \geq x_{n-k+2} = \dots = x_n \right\} \right\rangle,$$

where

$$A = \left\{ x = (x_1, \dots, x_n) \mid \sum_{i=1}^n x_i = hp, x_1, \dots, x_n \geq 0 \right\}.$$

Then K is a stable set if and only if

$$n \geq (h+1)(k-1).$$

Therefore, if n is large or k is small, then a minimal blocking coalition is formed and the rest of the players also form a coalition. Within each coalition, profits are shared equally. The next example deals with the case in which more than one types of raw material is necessary to produce a commodity.

Example 3 There are two types of raw materials P and Q, and one unit of each material is needed to produce one unit of an indivisible commodity, which is sold at p dollars. Player 1 holds one unit of raw material P, and each of players 2 and 3 holds one unit of raw material Q. This situation is formulated as the following characteristic function form game. The player set is $N = \{1, 2, 3\}$. Since one unit of raw materials P and Q is necessary to produce the commodity, the characteristic function v is given by

$$\begin{aligned} v(N) &= p, \\ v(\{1,2\}) &= v(\{1,3\}) = p, \quad v(\{2,3\}) = 0, \\ v(\{1\}) &= v(\{2\}) = v(\{3\}) = 0, \quad v(\emptyset) = 0. \end{aligned}$$

The set of imputations is

$$x = (x_1, x_2, x_3) \mid x_1 + x_2 + x_3 = p, x_1, x_2, x_3 \geq 0$$

The following set K is one of the stable sets in this game:

$$K = \{x = (x_1, x_2, x_3) \in A \mid x_2 = x_3\}.$$

To show internal stability, take two imputations $x = (x_1, x_2, x_3)$ and $y = (y_1, y_2, y_3)$ in K and suppose x dominates y . Then, the domination must hold via coalitions $\{1, 2\}$ and $\{1, 3\}$ since the value of the other coalitions except $\{1, 2, 3\}$ is 0. If $x \text{ dom}_{\{1,2\}} y$, then $x_1 > y_1$ and $x_2 > y_2$ hold. Thus, we have a contradiction $p = \sum_{i=1}^3 x_i > \sum_{i=1}^3 y_i = p$. The domination via $\{1, 3\}$ leads to the same contradiction. To show external stability, take any imputation $z = (z_1, z_2, z_3) \notin K$. Then $z_2 \neq z_3$. Without loss of generality, let $z_2 < z_3$. Define $y = (y_1, y_2, y_3)$ by

$$y_i = \begin{cases} z_1 + \frac{z_3 - z_2}{3} & \text{for } i = 1, \\ z_2 + \frac{z_3 - z_2}{3} & \text{for } i = 2, \\ z_2 + \frac{z_3 - z_2}{3} & \text{for } i = 3. \end{cases}$$

Then $y \in K$ and $y \text{ dom}_{\{1,2\}} z$, since $y_1 > z_1$, $y_2 > z_2$ and $y_1 + y_2 < v(\{1, 2\})$. This stable set shows that in negotiating on how to share the profit p dollars, players 2 and 3 form a coalition against player 1 and share equally the gain obtained through collaboration. There exist other stable sets in which players 2 and 3 collaborate, but they do not share equally the profit. More precisely, the following set

$$K = \{x = (x_1, x_2, x_3) \in A \mid x_2 \text{ and } x_3 \text{ move towards the same direction}\}$$

is a stable set, where the phrase “move towards the same direction” means that if x_2 increases then x_3 increases and if x_2 decreases then x_3 decreases.

A generalization of the above results is given by the following theorem due to Shapley (1959). Shapley’s original theorem is more complicated and holds in more general markets.

Theorem 6 Suppose there are m players, $1, 2, \dots, m$, each holding one unit of raw material P, and n players, $m+1, m+2, \dots, m+n$, each holding one unit of raw material Q. To produce one unit of an indivisible commodity, one unit of each of raw materials P and Q is necessary. One unit of commodity is sold at p dollars. In this market, the following set K is a stable set:

$$K = \{x = (x_1, x_2, \dots, x_{m+n}) \in A \mid x_1 = \dots = x_m, x_{m+1} = \dots = x_{m+n}\}.$$

where

$$A = \left\{ x = (x_1, \dots, x_m, x_{m+1}, \dots, x_{m+n}) \mid \sum_{i=1}^{m+n} x_i = p \times \min(m, n), x_1, \dots, x_{m+n} \geq 0 \right\},$$

is the set of imputations of this game.

This theorem shows that players holding the same raw material form a coalition and share equally the profit gained through collaboration. For further results on stable sets in production market games, refer to Hart (1973), Muto (1982a), and Owen (1995). Refer also to Lucas (1990), Owen (1968), and Shapley (1953) for further general studies on stable sets.

Assignment Games

The following three sections deal with applications of stable sets to two-sided markets including matching situations and barter markets with indivisible commodities. First, we consider the assignment market game introduced by Shapley and Shubik (1972). An assignment market consists of a set of $n_b(\geq 1)$ buyers $B = \{1, \dots, n_b\}$ and a set of $n'_s(\geq 1)$ sellers $F = \{1', \dots, n'_s\}$. Each seller $k' \in F$ has one indivisible unit of a commodity to sell, which we call object k' . Thus, we have n'_s objects in the market, and these objects may be differentiated. Each seller k' places a monetary value $c_{k'}$ (≥ 0) for object k' . Hereafter, we also denote by F the set of the n'_s objects. Each buyer $i \in B$ wants to buy at most one object in F and places a monetary value $h_{ik'}$ (≥ 0) for each object $k' \in F$.

If object k' is sold to buyer i , we have a surplus $u_{ik'} := \max(0, h_{ik'} - c_{k'})$. Let $c = (c_{k'})_{k' \in F}$, $H = (h_{ik'})_{(i,k') \in B \times F}$, and $U = (u_{ik'})_{(i,k') \in B \times F}$. An assignment market M is defined by the five elements (B, F, H, c, U) , where we will suppress H, c , or U when no confusion may arise. We remark that an assignment market with $|B| \neq |F|$ can be transformed into a market with $|B| = |F|$ by adding dummy buyers or sellers and zero rows or columns correspondingly to the original valuation matrix U .

An assignment game G is a characteristic function form game associated with a given assignment market $M = (B, F, U)$. We define the player set of G to be $B \cup F$. To define the characteristic function v of G , we first consider the following assignment problems $P(S)$ for each coalition $S \subseteq B \cup F$ with $S \cap B \neq \emptyset$ and $S \cap F \neq \emptyset$:

$$\begin{aligned} P(S) : \bar{m}(S) &= \max_x \sum_{(i,k') \in (S \cap B) \times (S \cap F)} u_{ik'} x_{ik'} \\ \text{s.t. } \sum_{k' \in S \cap F} x_{ik'} &\leq 1 \text{ for all } i \in S \cap B, \\ \sum_{i \in S \cap B} x_{ik'} &\leq 1 \text{ for all } k' \in S \cap F, \\ x_{ik'} &\geq 0 \text{ for all } (i,k') \in (S \cap B) \times (S \cap F). \end{aligned}$$

Each assignment problem $P(S)$ has at least one optimal integer solution (see Simonnard (1966)), which gives an optimal matching between sellers and buyers in S that gives the highest possible surplus in S . The characteristic function v is defined by $v(S) = \bar{m}(S)$ for each $S \subseteq B \cup F$ with $S \cap B \neq \emptyset$ and $S \cap F \neq \emptyset$, and $v(S) = 0$ for each S with $S \subseteq B$ or $S \subseteq F$, where the latter part means that the worth of a coalition consisting of only sellers or only buyers is zero, since those players cannot obtain any surplus by trading among themselves. We define $v(\emptyset) = 0$. The set of imputations is defined as the set

$$A = \left\{ (y, z) \in \mathcal{R}_+^B \times \mathcal{R}_+^F \mid \sum_{i \in B} y_i + \sum_{k' \in F} z_{k'} = v(B \cup F) \right\}.$$

Shapley and Shubik (1972) proved that for any assignment game G , the core C is nonempty and

given by the set of optimal solutions to the dual problem of assignment problem $P(B \cup F)$, i.e.,

$$C = \{(y, z) \in A \mid y_i + z_{k'} \geq u_{ik'} = v(\{i, k'\}) \text{ for each } (i, k') \in B \times F\}.$$

They also showed that for each $(y, z) \in C$, the vector $(z_{k'} + c_{k'})_{k' \in F}$ gives prices of the objects which equilibrate demand and supply of each object. A prototype of the assignment game was studied in detail by von Neumann and Morgenstern (1953). They considered stable sets of a market with two buyers and one seller having one object for sale.

Example 4 There are three players, seller 1' and buyers 1 and 2. Seller 1' has an object Q for sale. Seller 1' has no monetary value for Q. Buyers 1 and 2 value Q at h_1 and h_2 dollars, respectively. We assume $h_1 \geq h_2 \geq 0$ and $h_1 > 0$. The object Q can be owned by only one buyer, since it is indivisible. However, payments can be freely made among players.

Let G_1 be the assignment game of the above market. The player sets of G_1 are $B = \{1, 2\}$ and $F = \{1'\}$, and the characteristic function v is such that $v(\{1, 2, 1'\}) = h_1$, $v(\{1, 1'\}) = h_1$, $v(\{2, 1'\}) = h_2$, $v(\{1, 2\}) = 0$, and $v(1) = v(2) = v(1') = v(\emptyset) = 0$. The set of imputations is defined by $A = \{(y_1, y_2, z) \in \mathcal{R}_+^B \times \mathcal{R}_+^F \mid y_1 + y_2 + z = h_1\}$. We note that a payoff z of seller 1' also denotes a price of Q. To describe all stable sets in G_1 , let $\mathcal{F}$ be the set of continuous nondecreasing function f of $w \in [0, h_2]$ with $w \geq f(w) \geq 0$ and $f(0) = 0$. Von Neumann and Morgenstern (1953) showed that K is a stable set of G_1 if and only if K is a union of two subsets of imputations,

$$\begin{aligned} K_1 &= \{(y_1, y_2, z) \in A \mid y_1 = h_1 - z, y_2 = 0, h_1 \geq z \geq h_2\}, \\ K_2 &= \{(y_1, y_2, z) \in A \mid y_1 = h_1 - z - f(h_2 - z), \\ &\quad y_2 = f(h_2 - z), h_2 \geq z \geq 0\} \end{aligned}$$

supported by some $f \in \mathcal{F}$. The set K_1 is the core of G_1 . However, no imputation in K_1 dominates any imputation in $A_1 = \{(y_1, y_2, z) \in A \mid y_2 + z \leq h_2\}$. It is K_2 that gives K the full external stability

covering the area A_1 . Since K_2 is depicted as a curve in the area, it is called a bargaining curve. Since $\mathcal{F}$ contains infinitely many f , we have infinitely many stable sets in G_1 .

Let us consider implications of the stable sets in G_1 . We assume $h_1 > h_2$ for simplicity. Since a stable set K is a subset of imputations, we see that an efficient trade is made, i.e., Q is always assigned to buyer 1. More precisely, K_1 shows the possibility that buyer 2 with a lower reservation price h_2 is excluded by price competition, and buyer 1 pays $z \in [h_2, h_1]$ for Q. The set K_2 shows the possibility that the two buyers form a coalition to bargain with seller 1'. In this case, buyer 1 pays $z \in [0, h_2]$ for Q and gives buyer 2 part of the additional profit $h_2 - z$ as a side payment, which is determined by a function f . A stable set does not specify a particular price z . Instead, a stable set specifies a division rule of a profit. For G_1 , if a price z is determined in $[h_2, h_1]$, buyer 1 receives the full profit $h_1 - z$; if $z \in [0, h_2]$, a rule f specifies a portion of the additional profit $h_2 - z$ that buyer 2 gets. In addition, multiple division rules are allowed within $\mathcal{F}$, i.e., the set derived from internal and external stability. Interpreting these features, von Neumann and Morgenstern (1953) explained that a division rule in a stable set shows a standard of behavior that can be established among the players and that multiple division rules can be stable standards of behavior.

Example 4 raises two questions on the existence of stable sets in assignment games. First, when does an assignment game have the stable core? This is a natural question, since if $h_2 = 0$, then the set $\hat{K} = \{(y_1, y_2, z) \in A \mid y_1 = h_1 - z, y_2 = 0, h_1 \geq z \geq 0\}$ becomes the stable core. However, even if $h_2 > 0$, the set $\hat{K}$ is qualified as a stable set with no payment to buyer 2. This raises the second question: does every assignment game have a stable set in which any monetary transfer is restricted within each of the efficient trading pairs? This question was answered completely by recent studies.

To consider the first question, let $M = (B, F, U)$ be any assignment market and G the associated assignment game. We may assume without loss of generality that $|B| = |F| = n$ and that the rows and columns of U are arranged so that the diagonal assignment x^* with $x_{ii'}^* = 1 (i = 1, \dots, n)$ is an

optimal solution to $P(B \cup F)$. We say that U has a dominant diagonal if all of its diagonal entries are row and column maximums, i.e., $u_{ii'} = \max\{u_{ik'} | k' \in F\} = \max\{u_{ji'} | j \in B\}$ for each $i = 1, \dots, n$. If U has a dominant diagonal, each pair of buyer i and seller i' ($i = 1, \dots, n$) can produce a maximum surplus by trading with each other. Thus, each player does not have to compete with others in the same side for a partner. It is then proved that the dominant diagonal condition is a necessary and sufficient condition for the core of G to contain the imputations $(\bar{y}, \bar{z})$ and $(\bar{y}, \bar{z})$ with $\bar{y}_i = 0$, $\bar{z}_{i'} = u_{ii'}$, $\bar{y}_i = u_{ii'}$, and $\bar{z}_{i'} = 0$ for $i = 1, \dots, n$. Furthermore, the following theorem holds.

Theorem 7 Let $M = (B, F, U)$ be any assignment market with $|B| = |F|$. The associated assignment game G has the stable core if and only if U has a dominant diagonal.

It is also proved that an assignment game G is convex if and only if a valuation matrix U satisfies that $u_{ik'} = 0$ for each $(i, k') \in B \times F$ with $i \neq k$. The dominant diagonal condition is weaker than this condition for a convex assignment game. Hence, the core becomes the unique stable set for an assignment game in a class that includes convex assignment games. For more details, refer to Solymosi and Raghavan (2001).

To address the second question, we first show an instance of the Böhm-Bawerk market. A Böhm-Bawerk market is an assignment market (B, F, H, c) in which there is no product differentiation in objects for sale, i.e., $h_{ik'} = h_i$ for each $(i, k') \in B \times F$.

Example 5 There are two buyers 1 and 2 and two sellers 1' and 2'. Each seller has the same object Q . Sellers 1' and 2' value Q at $c_{1'}$ and $c_{2'}$ dollars, respectively, while buyers 1 and 2 value Q at h_1 and h_2 dollars, respectively. We assume $h_1 > h_2 > c_{2'} > c_{1'} \geq 0$.

Let G_2 be the associated assignment game. The set of imputations of G_2 is

$$A = \{(y, z) = (y_1, y_2, z_{1'}, z_{2'}) \in \mathcal{R}_+^4 \mid y_1 + y_2 + z_{1'} + z_{2'} = h_1 + h_2 - c_{1'} - c_{2'}\}$$

and the core C is

$$C = \{(y, z) \in A \mid (y_1, z_{1'}) = (h_1 - p_0, p_0 - c_{1'}), (y_2, z_{2'}) = (h_2 - p_0, p_0 - c_{2'}), p_0 \in [c_{2'}, h_2]\}.$$

Shubik (1985) presented an interesting transaction rule for Böhm-Bawerk markets, which is called the official price mechanism.

The official price mechanism presumes that a set of efficient trading pairs is determined and transactions are conducted only within each of those pairs. In this mechanism, an official price p is first announced, and then each pair trades at p . However, if the official price is not in between their reservation prices, the trade is done at the nearest reservation price. For Example 5, assume $(1, 1')$ and $(2, 2')$ to be the chosen efficient trading pairs. Then the official price mechanism derives the following set: $K_o = C_1 \cup C_2 \cup C \cup C_3 \cup C_4$, where

$$\begin{aligned} C_1 &= \{(y, z) \in A \mid (y_1, z_{1'}) = (0, h_1 - c_{1'}), (y_2, z_{2'}) = (0, h_2 - c_{2'})\} \text{ (obtained at } p \geq h_1) \\ C_2 &= \{(y, z) \in A \mid (y_1, z_{1'}) = (h_1 - p, p - c_{1'}), (y_2, z_{2'}) = (0, h_2 - c_{2'}), p \in [h_2, h_1]\}, \\ C &= \{(y, z) \in A \mid (y_1, z_{1'}) = (h_1 - p, p - c_{1'}), (y_2, z_{2'}) = (h_2 - p, p - c_{2'}), p \in [c_{2'}, h_2]\}, \\ C_4 &= \{(y, z) \in A \mid (y_1, z_{1'}) = (h_1 - p, p - c_{1'}), (y_2, z_{2'}) = (h_2 - c_{2'}, 0), p \in [c_{1'}, c_{2'}]\}, \\ C_5 &= \{(y, z) \in A \mid (y_1, z_{1'}) = (h_1 - c_{1'}, 0), (y_2, z_{2'}) = (h_2 - c_{2'}, 0)\} \text{ (obtained at } p \leq c_{1'}) \end{aligned}$$

The set K_o contains only imputations brought by transactions within each of the efficient trading pairs $(1, 1')$ and $(2, 2')$. There is no payment to a third

party in the sense that there is no payment between players belonging to different trading pairs. Shubik (1985) proved the following theorem.

Theorem 8 For any Böhm-Bawerk market, the official price mechanism gives a stable set that consists of only imputations brought by transactions within each pair belonging to a given set of efficient trading pairs.

Example 5 further suggests that a stable set with no payment to a third party can be regarded as a union of cores of subgames defined below. For a given assignment market $M = (B, F, U)$, let x^* be any optimal solution to the associated assignment problem $P(B \cup F)$. A buyer i is said to be matched at x^* if $x_{ik'}^* = 1$ for some $k' \in F$. A matched seller is defined in the same way. Let B^*, F^* be the sets of matched buyers and sellers at x^* . A bijection μ from B^* onto F^* is referred to as the optimal matching associated with x^* if $\mu(i) = j'$ and $\mu^{-1}(j') = i$ for $(i, j') \in B^* \times F^*$ with $x_{ij'}^* = 1$. We simply call μ an optimal matching in M if μ is an optimal matching associated with some optimal solution to $P(B \cup F)$. For any pair of subsets $I \subset B$ and $J \subset F$, a submarket $M_{-I \cup J}$ is an assignment market $(B \setminus I, F \setminus J, U_{-I \cup J})$, where $U_{-I \cup J}$ is the valuation submatrix obtained by deleting the rows in I and the columns in J from U . The subgame $G_{-I \cup J}$ is the assignment game associated with $M_{-I \cup J}$. The characteristic function of $G_{-I \cup J}$ is denoted by $v_{-I \cup J}$. Given an optimal matching μ in M , a subgame $G_{-I \cup J}$ of G is said to be μ -compatible if

$$v(B \cup F) = v_{-I \cup J}((B \setminus I) \cup (F \setminus J)) + \sum_{i \in I \cap B^*} u_{i\mu(i)} + \sum_{k' \in J \cap F^*} u_{\mu^{-1}(k')k'}$$

If $G_{-I \cup J}$ is a μ -compatible subgame of G , the restriction of μ defined on $B^* \setminus (I \cup \mu^{-1}(J \cap F^*))$ gives an optimal matching in the submarket $M_{-I \cup J}$. Note that G is also a μ -compatible subgame with $I = J = \emptyset$. Let S^μ be the set of pairs $(I, J) \in 2^B \times 2^F$ whose deletion gives a μ -compatible subgame of G . For each $(I, J) \in S^\mu$, let $C_{-I \cup J}$ be the core of $G_{-I \cup J}$. We then define the extended core $\bar{C}_{-I \cup J}$ to be a subset of the imputation set A of G with the property that $(y_{B \setminus I}, y_I, z_{F \setminus J}, z_J) \in \bar{C}_{-I \cup J}$ if and only if:

1. $(y_{B \setminus I}, z_{F \setminus J}) \in C_{-I \cup J}$
2. $y_i = u_{i\mu(i)}$ for $i \in I \cap B^*$
3. $z_{k'} = u_{\mu^{-1}(k')k'}$ for $k' \in J \cap F^*$
4. $y_i = z_{k'} = 0$ for $i \in I \setminus B^*$ and $k' \in J \setminus F^*$

Applying the above notions, the stable set K_O of Example 5 is expressed as follows:

$$K_O = \bar{C}_{-\{1', 2'\}} \cup \bar{C}_{-\{2'\}} \cup \bar{C}_{-\emptyset} \cup \bar{C}_{-\{2\}} \cup \bar{C}_{-\{1, 2\}}$$

In general, we have the following theorem, whose proof was outlined by Shubik (1985) and completed by Núñez and Rafels (2013).

Theorem 9 Let $M = (B, F, U)$ be any assignment market and G its associated assignment game. For each optimal matching μ in M , the union of the extended cores of μ -compatible subgames, i.e., $K_\mu = \bigcup_{(I, J) \in S^\mu} \bar{C}_{-I \cup J}$ gives a stable set of G . Furthermore, K_μ is a unique stable set that we have when monetary transfer is restricted within each pair formed at μ .

The existence of stable sets in the whole class of assignment games had been an unresolved question for many years, which was positively answered by Theorem 9.

Marriage Games

This section considers stable sets of one-to-one matching games, the so-called marriage games. A marriage game G is defined by a triple (M, W, R) . The elements M and W are two disjoint finite sets of players: $M = \{m_1, \dots, m_b\}$ and $W = \{w_1, \dots, w_g\}$. Each player in one set wants to form a pair with a player in the other set. We may interpret M and W as sets of men and women, and therefore we call this game a marriage game. The element R denotes a preference profile $(R_i)_{i \in M \cup W}$. We mean by $yR_i z$ that $yP_i z$ or $yI_i z$ holds, where $yP_i z$ means that player i strictly prefers y to z and $yI_i z$ means that i is indifferent with y and z . We assume that each $m \in M$ has strict preferences over $W \cup \{m\}$, namely, (1) $xP_m y$ or $yP_m x$ for each $x, y \in W \cup \{m\}$ with $x \neq y$ and (2) $xI_m y$ if and only if $x = y$. We also assume that each $w \in W$ has strict preferences over $M \cup \{m\}$. For each player $i \in M \cup W$, the choice

i means staying single. A potential partner who is preferred (or inferior) to i is said to be acceptable (or unacceptable) for player i . For simplicity, we also use a rank order list to present a preference relation. For example, “ $(w)w_1, \dots, w_k, m, w_{k+1}, \dots, w_g$ ” is the rank order list representing m ’s preferences $w_1 P_m \dots P_m w_k P_m m P_m w_{k+1} P_m \dots P_m w_g$.

An outcome of a marriage game $G = (M, W, R)$ is a matching, which is defined by a bijection $\mu: M \cup W \rightarrow M \cup W$ with $\mu(m) \in W \cup \{m\}$ for each $m \in M$, $\mu(w) \in M \cup \{w\}$ for each $w \in W$, and $\mu(m) = w$ if and only if $\mu(w) = m$. Let A_0 be the set of all matchings in G . Given a matching $\mu \in A_0$, a player $i \in M \cup W$ is said to be matched at μ if μ assigns i to another player in the other set. Otherwise, player i is said to be unmatched at μ . A pair $(m, w) \in M \times W$ is called a matched pair at μ if $\mu(m) = w$ (and $\mu(w) = m$). We denote by M_μ and W_μ the sets of matched men and women at μ , respectively. For simplicity, we also use set theoretic notations of a matching such as $\mu = \{(m_i, w_{j_1}), \dots, (m_{i_k}, w_{j_k})\}$ and $(m_i, w_{j_1}) \in \mu$, where $(m_i, w_{j_1}), \dots, (m_{i_k}, w_{j_k})$ are the matched pairs at μ .

To define the core and a stable set of a marriage game G , we define a domination relation between two matchings. Let μ, ν be any pair of matchings. We say that μ dominates ν if (1) there exists a matched pair (m, w) at μ with $\mu(m) = w P_m \nu(m)$ and $\mu(w) = m P_w \nu(w)$ or (2) there exists an unmatched player $a \in M \cup W$ at μ with $\mu(a) = a P_a \nu(a)$. When μ dominates ν via a matched pair (m, w) at μ , the pair (m, w) is called a blocking pair to ν . If a coalition $S \subset M \cup W$ is effective in μ ’s domination of ν , i.e., $\mu(S) = S$ and $\mu(a) P_a \nu(a)$ for each $a \in S$, then S always includes a matched pair or an unmatched player at μ that enables μ to dominate ν . Thus, we only consider domination by a pair or a single player. When we do not have to specify a matching that contains a blocking pair (m, w) , we say that ν is blocked by (m, w) . A matching ν can also be blocked by a single player $i \in M \cup W$ if $i P_i \nu(i)$. A matching μ is said to be individually rational if μ is not blocked by any single player. Let A be the set of individually rational matchings of G . Finally, we say that μ is M -Pareto superior to ν if $\mu(m) R_m \nu(m)$ for each $m \in M$ with strict preference for some m , and μ is W -Pareto superior to ν if $\mu(w) R_w \nu(w)$ for each $w \in W$ with strict preference for some w . The μ ’s

Pareto-inferiority to ν for M and for W is defined in the same way by replacing $\mu(m) R_m \nu(m)$ with $\nu(m) R_m \mu(m)$ and $\mu(w) R_w \nu(w)$ with $\nu(w) R_w \mu(w)$.

The core C of a marriage game G is defined to be a subset of individually rational matchings that are not blocked by any pair $(m, w) \in M \times W$. Although a matching in the core is usually called a stable matching in the literature on matching games, we call it a core matching to avoid confusion with matchings in a stable set. Gale and Shapley (1962) proved that the core is nonempty for any marriage game by using their celebrated deferred acceptance algorithm.

Let V be any nonempty set of matchings. For any pair of matchings $\mu, \nu \in V$, we define two functions $\mu \wedge \nu$ and $\mu \vee \nu$ from $M \times W$ to $M \times W$ by

$$\begin{aligned} \mu \wedge \nu(m) &= \min\{\mu(m), \nu(m)\} \text{ for each } m \in M, \\ \mu \wedge \nu(w) &= \max\{\mu(w), \nu(w)\} \text{ for each } w \in W, \\ \mu \vee \nu(m) &= \max\{\mu(m), \nu(m)\} \text{ for each } m \in M, \\ \mu \vee \nu(w) &= \min\{\mu(w), \nu(w)\} \text{ for each } w \in W, \end{aligned}$$

where $\min\{\mu(i), \nu(i)\}$ and $\max\{\mu(i), \nu(i)\}$, respectively, denote a weakly inferior element and a weakly preferable element in $\{\mu(i), \nu(i)\}$ for player $i \in M \cup W$. We say that V is a lattice if $\mu \wedge \nu \in V$ and $\mu \vee \nu \in V$ for each $\mu, \nu \in V$. We also say that V has invariant matched players if $M_\mu = M_\nu$ and $W_\mu = W_\nu$ for each $\mu, \nu \in V$.

It is well known that the core C of a marriage game has the following properties:

1. Lattice structure: C is a lattice.
2. Invariant matched players: C has invariant matched players.
3. Opposition of interests: for any $\mu, \nu \in C$, if $\mu(m) R_m \nu(m)$ for each $m \in M$, then $\nu(w) R_w \mu(w)$ for each $w \in W$ and vice versa.
4. Existence of polarized optimal core matchings: there exists a man-optimal core matching μ_M and a woman-optimal core matching μ_W such that μ_M and μ_W are M -Pareto superior and W -Pareto superior to any other $\mu \in C$, respectively. Here, if C is a singleton, these two matchings coincide and vice versa.

The third and fourth properties are in fact derived by the lattice property. The polarized

optimal core matchings can be found by the deferred acceptance algorithm in polynomial time. See Roth and Sotomayor (1990), Gusfield and Irving (1989), and Manlove (2013) for more details of the core and the deferred acceptance algorithm.

A stable set of a marriage game is a nonempty subset K of A (the set of individually rational matchings) that satisfies internal stability (i.e., for any $\mu, \nu \in K$, μ does not dominate ν) and external stability (i.e., for each $\nu \in A \setminus K$, there exists $\mu \in K$ that dominates ν). We use $C(R)$ and $\mathcal{K}(R)$ to denote the core and the set of stable sets obtained under a preference profile R . For a stable set defined on A_0 (the set of all matchings), see the remark at the end of this section. Let us examine a simple marriage game to see its core and stable set.

Example 6 Let $G_3 = (M, W, R)$ be the marriage game with $M = \{m_1, m_2, m_3\}$, $W = \{w_1, w_2, w_3\}$, and the following preference profile R :

$$\begin{array}{ll} m_1) w_2, w_1, w_3, m_1, & w_1) m_3, m_2, m_1, w_1 \\ m_2) w_2, w_3, w_1, m_2, & w_2) m_3, m_1, m_2, w_2 \\ m_3) w_3, w_2, w_1, m_3, & w_3) m_1, m_2, m_3, w_3. \end{array}$$

Using the deferred acceptance algorithm, we find the man-optimal core matching $\mu_M = \{(m_1, w_1), (m_2, w_3), (m_3, w_2)\}$ and the woman-optimal core matching $\mu_W = \{(m_1, w_3), (m_2, w_1), (m_3, w_2)\}$. Except for μ_M and μ_W , there is no other matching μ with $\mu_M(m)R_m \mu(m)R_m \mu_W(m)$ for each $m \in M$. Thus, we have $C(R) = \{\mu_M, \mu_W\}$ from the lattice property of the core.

We see that each matching with unmatched players is dominated by a core matching. However, the core $C(R)$ is not externally stable, since the matching $\mu_M^* = \{(m_1, w_1), (m_2, w_2), (m_3, w_3)\}$ is not dominated by any core matching. On the other hand, since each core matching is not dominated by any matching, the set $K = \{\mu_M^*, \mu_M, \mu_W\}$ is internally stable. We show the external stability of K . Although μ_M^* is dominated by $\mu_1 = \{(m_1, w_2), (m_2, w_1), (m_3, w_3)\}$ via (m_1, w_2) , μ_1 is dominated by $\mu_M \in K$ via (m_2, w_3) . The remaining two matchings $\mu_2 = \{(m_1, w_2),$

$(m_2, w_3), (m_3, w_1)\}$ and $\mu_3 = \{(m_1, w_3), (m_2, w_2), (m_3, w_1)\}$ are dominated by μ_W via (m_3, w_2) . Thus, K is externally stable and therefore a stable set of G_3 . We note that K has a lattice structure and invariant matched players. Furthermore, the set K is in fact the unique stable set of G_3 .

Ehlers (2007) gave the following two characterizations of a stable set in a marriage game with a remark that the second type of characterization was first noted by von Neumann and Morgenstern (1953).

Theorem 10 Let $G = (M, W, R)$ be a marriage game.

1. If K is a stable set of G , then K is a maximal set such that:
 - (a) K is a superset of $C(R)$.
 - (b) K is a lattice.
 - (c) K has invariant matched players.

Furthermore, if K is a unique maximal set with (a), (b), and (c), then K is a stable set of G .

2. K is a stable set of G if and only if

$$K = \{\mu \in A \mid \mu \text{ is not blocked by any pair in } T(K)\}.$$

where $T(K)$ is the set of matched pairs at some matching in K , i.e.,

$$T(K) = \{(m, w) \in M \times W \mid (m, w) \in \mu \text{ for some } \mu \in K\}.$$

The first characterization has interesting implications. First, a stable set has polarized optimal matchings from the lattice property. In Example 6, μ_M^* and μ_W are man-optimal and woman-optimal in K , respectively. Here, μ_W is also the woman-optimal core matching. Second, since a stable set is a superset of the core, a man (woman)-optimal matching in a stable set can be M -Pareto (W -Pareto) superior to a man (woman)-optimal core matching, respectively. Third, a stable set has the same invariant matched players as the core. Thus, if a player is unmatched at a core matching, this player is also unmatched at any matching in any stable set. Finally, this characterization gives a sufficient condition for the existence of a stable set, which is the unique maximal

set satisfying conditions (a), (b), and (c). Ehlers (2007) gave an example in which we have two maximal sets with (a), (b), and (c), but only one of them is a stable set. We will discuss the existence of a stable set later. The second characterization means that a stable set K is equivalent to the set of individually rational matchings that are not dominated by any pair in $T(K)$. Thus, a stable set in a marriage game can be regarded as a core obtained when only pairs in $T(K)$ can block a matching.

For any preference profile R and any set of pairs $S \subset M \times W$, we define the preference profile $R \setminus S$ to be a preference profile that is obtained by the following trimming-off operation on the rank order lists corresponding to R :

- For each $m \in M$, on the rank order list of R_m , remove all acceptable women w with $(m, w) \in S$, and put them immediately below m without changing their relative ranks.
- For each $w \in W$, on the rank order list of R_w , remove all acceptable men m with $(m, w) \in S$, and put them immediately below w without changing their relative ranks.

Given a stable set K under R , let

$$S = \{(m, w) \in M \times W \mid (m, w) \notin \mu \text{ for each } \mu \in K\} \\ (= (M \times W) \setminus T(K))$$

and $\bar{R} = R \setminus S$. We call $\bar{R}$ the K -trimmed preference profile of R . Then the second characterization means the following equivalence:

$$K = \{\mu \in \mathcal{A} \mid \mu \text{ is not blocked by any } (m, w) \in T(K)\} \\ = C(\bar{R}).$$

In Example 6, since $(m_1, w_2), (m_3, w_1) \notin T(K)$, we obtain the K -trimmed preference profile $\bar{R}$ by trimming off these two pairs from R as shown below:

$$\begin{array}{ll} m_1 \left. \begin{array}{l} w_1, w_3, m_1, w_2, \\ m_2 \end{array} \right\} w_1 \left. \begin{array}{l} m_2, m_1, w_1, m_3 \\ m_3 \end{array} \right\} m_2, m_2, w_2, m_1 \\ m_2 \left. \begin{array}{l} w_2, w_3, w_1, m_2, \\ m_3 \end{array} \right\} w_2 \left. \begin{array}{l} m_3, m_2, w_2, m_1 \\ m_1 \end{array} \right\} m_1, m_2, m_3, w_3. \\ m_3 \left. \begin{array}{l} w_3, w_2, m_3, w_1 \\ \end{array} \right\} w_3 \left. \begin{array}{l} m_1, m_2, m_3, w_3. \end{array} \right\} \end{array}$$

For example, m_1 's rank order list is obtained by removing w_2 from the original list w_2, w_1, w_3, m_1

and putting it immediately below m_1 . The original core $C(R) = \{\mu_M, \mu_W\}$ is extended to the core $C(\bar{R}) = \{\mu_M^*, \mu_M, \mu_W\}$ under $\bar{R}$, and $C(\bar{R})$ gives the stable set K under R . We can easily make the K -trimmed preference profile if a stable set K is given. However, without knowing a stable set K , can we make a preference profile R under which the core $C(R)$ gives a stable set K under R ? The following theorem gives a positive answer to this question.

Theorem 11 For any marriage game $G = (M, W, R)$, there exists a preference profile R^* with the property that $\mathcal{K}(R) = \mathcal{K}(R^*) = \{C(R^*)\}$. The preference profile R^* can be constructed from R in polynomial time.

The equation $\mathcal{K}(R) = \mathcal{K}(R^*) = \{C(R^*)\}$ means that the set of stable sets under R^* is equivalent to the set of stable sets under the original preference profile R and that a marriage game (M, W, R) has a unique stable set, which is $C(R^*)$. As mentioned in section "Stable Sets in Characteristics Function from Games," the core existence property of marriage games does not suffice for the existence of a stable set. Since the preference profile R^* can be obtained for any marriage game, this theorem shows the existence of a stable set in any marriage game.

Outline of the proof. The above theorem is proved by using interesting properties of cores and stable sets of marriage games. In the following, we show an outline of the proof that was given by Wako (2010):

1. **Preliminaries:** Since any stable set includes the core and has invariant matched players from Theorem 10, if some players are unmatched at some core matching, then they are also unmatched at any matching in any stable set. Thus, for finding a stable set, it suffices to consider a subgame (M', W', R') of G in which (1) M' and W' do not include the unmatched players in G and (2) $R' = (R'_x)_{x \in M \cup W}$ is such that each R'_m over $W' \cup \{m\}$ is defined by using the relative rankings over $W' \cup \{m\}$ under R_m , and each R'_w over $M' \cup \{w\}$ is also defined in the same way.

Furthermore, since we are considering a stable set defined on the set of individually rational matchings, we can neglect a matching having a pair in which one player is unacceptable to the other. Thus, we may delete such pairs by trimming them off from the preference profile. Hence, without loss of generality, we may assume that all players in $M \cup W$ are matched at some core matching, and for each $(m, w) \in M \times W$, w is acceptable to m if and only if m is acceptable to w .

Under the above assumption, take any preference profile R and any core matching $\mu \in C(R)$. Let $A(R)$ be the set of individually rational matchings under R . It should be noted that all players are matched at μ . First, we define an unstable pair under R to be any pair $(m, w) \in M \times W$ that is not formed at any matching in any stable set $K \in \mathcal{K}(R)$. We say that:

- $v \in A(R)$ is a W -inferior μ -adjacent matching under R if:
 1. All players are matched at v .
 2. For each $w \in W$, $v(w)$ is ranked immediately below $\mu(w)$ in R_w or $v(w) = \mu(w)$.
 3. v is M -Pareto superior to μ .
- (m, w) is a W -inferior μ -adjacent pair under R if w and m are mutually acceptable, and m is ranked immediately below $\mu(w)$ in R_w .
- $v \in A(R)$ is a W -worst matching under R if for each $w \in W$, $v(w)$ is the least preferred acceptable partner R_w .

We also define an M -inferior μ -adjacent matching, an M -inferior μ -adjacent pair, and an M -worst matching in the same manner as above by exchanging the set W and its element w for the set M and its element m symmetrically.

We note that even if μ is fixed, we have different W (or M)-inferior μ -adjacent matchings and pairs depending on R . It is well known that if μ is a core matching in $C(R)$, then both W -inferior μ -adjacent and M -inferior μ -adjacent matchings also belong to $C(R)$. In addition, these matchings can be found by the method of elimination of rotations. However, if μ is W -worst (M -worst), then there exists no W -inferior (M -inferior) μ -adjacent matching. The same is true for

both the W -inferior and M -inferior μ -adjacent pairs. Refer to Gusfield and Irving (1989) for the details of these properties.

2. **Important properties:** Wako (2010) showed the following properties of core matchings, unstable pairs, and sets of stable sets.

Property 1 (1) If a core matching $\mu \in C(R)$ dominates each matching $v \in A(R)$ containing (m, w) , then (m, w) is an unstable pair under R .

(2) Let μ be a core matching in $C(R)$ that is neither W -worst nor M -worst.

(2a) If μ has no W -inferior μ -adjacent matching, then there exists a W -inferior μ -adjacent pair (m, w) such that μ dominates each matching $v \in A(R)$ containing (m, w) .

(2b) If μ has no M -inferior μ -adjacent matching, then there exists an M -inferior μ -adjacent pair (m, w) such that μ dominates each matching $v \in A(R)$ containing (m, w) .

Property 2 If (m, w) is an unstable pair under R , then $\mathcal{K}(R) = \mathcal{K}(R \setminus \{(m, w)\})$.

Property 1 means that for any core matching μ under a given preference profile R , if there is no W (M)-inferior μ -adjacent matching, then there exists an unstable pair, or it is the W -worst or M -worst core matching. Property 2 means that even if we trim off an unstable pair from R , the set of stable sets does not change. Using these properties, we will present a basic idea to prove Theorem 11.

3. **Procedure to find a stable set:** First, let $R^{(0)} = R$. Starting with the woman-optimal core matching $\mu \in C(R^{(0)})$, we examine whether there is a W -inferior μ -adjacent matching. If we find such a matching μ' , which belongs to $C(R^{(0)})$, then move on to μ' and examine whether there is W -inferior μ' -adjacent matching. If no such a matching exists, then from (1) and (2a) in Property 1, there is a W -inferior μ' -adjacent pair (m_1, w_1) that is unstable under $R^{(0)}$, and then trim it off from $R^{(0)}$ and let $R^{(1)} = R^{(0)} \setminus \{(m_1, w_1)\}$. From Property 2, we have $\mathcal{K}(R^{(0)}) = \mathcal{K}(R^{(1)})$. By this trimming-off operation, the core gets

larger, i.e., $C(R^{(0)}) \subseteq C(R^{(1)})$. Thus, $\mu' \in C(R^{(1)})$. Applying Property 1 to μ' and $R^{(1)}$, we examine whether there is a W -inferior μ' -adjacent matching under $R^{(1)}$. If no such a matching exists, then from (1) and (2a) in Property 1, there is a W -inferior μ' -adjacent pair (m_2, w_2) that is unstable under $R^{(1)}$, and then trim it off from $R^{(1)}$ and let $R^{(2)} := R^{(1)} \setminus \{(m_2, w_2)\}$. At this point, we have

$$\mathcal{K}(R^{(0)}) = \mathcal{K}(R^{(1)}) = \mathcal{K}(R^{(2)}) \text{ and} \\ C(R^{(0)}) \subseteq C(R^{(1)}) \subseteq C(R^{(2)}).$$

Repetition of the above operation keeps the same set of stable sets while the core is getting larger. Furthermore, the number of individually rational pairs (i.e., pairs of mutually acceptable partners) that can be formed under an existing preference profile becomes smaller one by one each time an unstable pair is trimmed off. In some k_1 th repetition of the trimming-off operation, we obtain a preference profile $R^{(k_1)}$ and a matching $\mu_M^* \in C(R^{(k_1)})$ such that μ_M^* is a W -worst matching under $R^{(k_1)}$. Then, starting with μ_M^* , we examine whether there is an M -inferior μ_M^* -adjacent matching. We repeat the same operation as above with respect to M -inferiority. In some k_2 -th repetition ($k_2 \geq k_1$), we obtain a preference profile $R^* (= R^{(k_2)})$ and a matching $\mu_W^* \in C(R^*)$ such that μ_W^* is an M -worst matching under R^* . At this point, the procedure ends, since we have neither an M -inferior μ_W^* -adjacent matching nor such a pair. We then see that each individually rational pair under R^* is formed at some core matching in $C(R^*)$ and that $C(R^*)$ is a stable set under R^* . If the core is a stable set, it is a unique stable set. Hence, we have

$$\mathcal{K}(R^{(0)}) = \dots = \mathcal{K}(R^*) = \{C(R^*)\}.$$

The matchings μ_M^* and μ_W^* are the man-optimal and woman-optimal matchings belonging to the stable set $C(R^*)$. Theorem 11 is in fact proved by using an algorithm that is based on the above outline. In the algorithm, every unstable pair is exhaustively found by applying the method of

elimination of rotations, which is studied in detail by Gusfield and Irving (1989). For the details of the proof of Theorem 11, see Wako (2010).

Remark In this section, we defined a stable set on the set A of individually rational matchings. However, it can also be defined on the set A_0 of all matchings. Ehlers (2007) defined the way of modifying a given preference profile so that we can have a stable set on A_0 as a stable set on the set of individually rational matchings under the modified preference profile. Hence, the theorems in this section hold for stable sets on the set of all matchings.

Bando (2014) considered a strictly strong Nash equilibria of a preference revelation marriage game G' with the women's preference profile R_W^0 fixed. In G' , when each man m reports his preference R_m , which may not give their true preference profile R_M^0 , the man-optimal core matching at (R_M, R_W^0) is given as an outcome. Bando proved that for any true preference profile $R^0 = (R_M^0, R_W^0)$, the man-optimal matching in the stable set K at R^0 is obtained as one of the possibly multiple strictly strong Nash outcomes in G' .

House Barter Games

In this section, we consider a market in which only indivisible commodities are bartered without using monetary transfers. This market was originally considered by Shapley and Scarf (1974). The market we consider has n (≥ 2) players, each endowed with one indivisible unit of a commodity, which we call a house. Let $N = \{1, \dots, n\}$ be the set of players. The n houses may be differentiated, and the house initially owned by player i is referred to as house i . We use N also to denote the set of houses in the market. Each player i has a complete, reflexive, and transitive preference relation R_i over N . The bundle $R = (R_i)_{i \in N}$ is a preference profile. We assume that each player wants to own exactly one house and strictly prefers owning a house to owning no house. Furthermore, there is no divisible good such as money in the market. The players only exchange their houses in a mutually beneficial way. We define

an outcome of the market, called an allocation, to be a bijection x from N onto N , where $x(i)$ denotes the house assigned to player i at x . An allocation is a permutation of N . For simplicity, we also use a vector representation such as $x = (x_1, \dots, x_n)$ to denote an allocation x , in which each element x_i denotes $x(i)$. Let A be the set of allocations. A market defined as above is referred to as a house barter market $M = (N, R)$ or briefly a market M .

Let x, y be any pair of allocations in a market $M = (N, R)$. For each nonempty coalition $S \subseteq N$, let $x(S)$ be the set of houses assigned to the members of S at x , i.e.,

$$x(S) = \{j \in N \mid j = x(i) \text{ for some } i \in S\}.$$

We say that x weakly dominates y and denote it by x wdom y if there exists a coalition S such that:

1. $x(i)R_i y(i)$ for each $i \in S$ with strict preference for some $i \in S$.
2. $x(S) = S$.

The second condition is the effectivity condition, which requires that each player i in S can obtain house $x(i)$ by exchanging their own endowments. We say that x strongly dominates y and denote it by x sdom y if $x(i)P_i y(i)$ for each $i \in S$ and $x(S) = S$. We use the notations x wdom _{S} y and x sdom _{S} y when we indicate the associated coalition S .

An allocation x is individually rational if $x(i)R_i i$ for each player $i \in N$. An allocation x is Pareto efficient if there exists no allocation $y \in A$ with y wdom _{N} x . If there is no $y \in A$ with y sdom _{N} x , then x is weakly Pareto efficient. The three sets of individually rational, Pareto efficient, and weakly Pareto efficient allocations are denoted by IR , PA , and WPA , respectively.

In this section, we consider the two types of cores and stable sets defined by the weak and strong notions of domination. The strict core (or strong core) of M is the set of allocations that are not weakly dominated by any other allocations. A wdom stable set of M is a nonempty set K of allocations with internal stability and external stability defined by weak domination, i.e., for any

$\mu, v \in K$, μ does not weakly dominate v , and for each $v \in A \setminus K$, there exists $\mu \in K$ that weakly dominates v . On the other hand, the core of M is the set of allocations that are not strongly dominated by any other allocations. An sdom stable set of M is a nonempty set of allocations with internal stability and external stability defined by strong domination instead of weak domination.

From the above definitions, the strict core is a subset of $PA \cap IR$, and the core is a subset of $WPA \cap IR$. A wdom stable set is a subset of PA , and an sdom stable set is a subset of WPA . However, both wdom and sdom stable sets may not be subsets of IR . Shapley and Scarf (1974) proved that the core is nonempty for all house barter markets. However, since external stability is not imposed on the core, the core does not necessarily coincide with an sdom stable set. In fact, the following example shows that there is a house barter market with no sdom stable set.

Example 7 Let $M_1 = (N, R)$ be the market with the player set $N = \{1, 2, 3\}$ and the following preference profile:

- 1) $2P_1 3P_1 1$,
- 2) $3P_2 1P_2 2$,
- 3) $1P_3 2P_3 3$.

Market M_1 has six allocations: $x^1 = (2, 3, 1)$, $x^2 = (2, 1, 3)$, $x^3 = (1, 3, 2)$, $x^4 = (3, 2, 1)$, $x^5 = (3, 1, 2)$, and $x^6 = (1, 2, 3)$. From the preference profile, x^1 is clearly a core allocation. Suppose that M_1 has an sdom stable set K . Then we must have $x^1 \in K$ from the external stability of K . Although x^1 strongly dominates x^5 and x^6 , it does not strongly dominate any $x^k \in \{x^2, x^3, x^4\}$ ($\equiv X$). Furthermore, we have x^2 sdom _{$\{1, 2\}$} x^4 , x^4 sdom _{$\{1, 3\}$} x^3 , and x^3 sdom _{$\{2, 3\}$} x^2 . From the internal stability of K , we see that only one allocation $x^k \in X$ can be contained in K . However, the allocation x^k can strongly dominate only one allocation in $X \setminus \{x^k\}$. This contradicts that K has external stability. Thus, no sdom stable set exists in M_1 .

The market M_1 , however, has the following features. The singleton $\{x^1\}$ is the strict core, and x^1 weakly dominates all the other allocations

$x^2, \dots, x^6$. In addition, every player shows only strict preferences. Roth and Postlewaite (1977) in fact proved that for any house barter market, if each player's preferences are strict, then the strict core is a singleton, and it is a unique wdom stable set, i.e., the wdom stable core. Wako (1991) proved that this property is generalized as follows:

Theorem 12 For any house barter market $M = (N, R)$, if the strict core SC is nonempty, it is a unique wdom stable set. Furthermore, for any $x, y \in SC$, we have $x(i)Iy(i)$ for each $i \in N$.

This theorem shows the following features of the strict core of a house barter market. First, any allocation outside the strict core is weakly dominated by some strict core allocation, since the strict core is a wdom stable set. Second, even if the strict core contains different allocations, they are indifferent for each player. However, the strict core can be empty when indifferences are allowed in preferences. Shapley and Scarf (1974) first pointed out this fact by the following example.

Example 8 Let $M_2 = (N, R)$ be the market with the player set $N = \{1, 2, 3\}$ and the following preference profile:

- 1) $2P_13I_11$,
- 2) $1I_23P_22$,
- 3) $2P_31I_33$.

We see that the strict core of M_2 is empty and that the sets $K^1 = \{(2, 3, 1), (2, 1, 3)\}$ and $K^2 = \{(1, 3, 2), (3, 1, 2)\}$ are both wdom stable sets of M_2 . Thus, neither the nonemptiness of the strict core nor the uniqueness of a wdom stable set holds when indifferences are allowed in preferences.

Quint and Wako (2004) considered a necessary and sufficient condition for the strict core to be nonempty. For each player $i \in N$ and each nonempty coalition $S \subseteq N$, let $B_i(S)$ be the set of player i 's most preferred houses in S , i.e., $B_i(S) = \{h \in S \mid hR_ji \text{ for each } j \in S\}$. We call a partition $T = \{T_1, \dots, T_m\}$ of N a partition by minimal self-mapped sets (PMSS) if each $T_k \in T$ satisfies the following conditions:

$$T_k = \bigcup_{i \in T_k} B_i(N \setminus \bigcup_{l=1}^{k-1} T_l) \text{ and there is no nonempty } S \subset T_k \text{ with } S = \bigcup_{i \in S} B_i(N \setminus \bigcup_{l=1}^{k-1} T_l).$$

We say that $T_k \in T$ is a lower (higher) set of $T_l \in T$ if $k > l$ ($k < l$). The fact that $T = \{T_1, \dots, T_m\}$ is a PMSS means that for each player i in $T_k \in T$, player i 's most preferred houses among T_k and its lower sets are endowed in T_k . Quint and Wako (2004) showed that any house barter market has at least one PMSS and that even if more than one PMSS exists, each PMSS consists of the same sets with only orders of some sets being different. Then the following theorem was proved.

Theorem 13 Let $T = \{T_1, \dots, T_m\}$ be a PMSS of a house barter market $M = (N, R)$. Then the strict core is nonempty if and only if there exists an allocation $x \in A$ such that

$$x(T_k) = T_k \text{ and } x(i) \in B_i(N \setminus \bigcup_{l=1}^{k-1} T_l) \text{ for each } i \in T_k \text{ and each } T_k \in T.$$

The necessary and sufficient condition given above requires that in each $T_k \in T$, each player $i \in T_k$ can obtain his/her most preferred house (among those owned in T_k and its lower sets) through a feasible exchange within T_k . We refer to this condition as segmentability. Suppose that T is a PMSS of a house barter market with segmentability. Then, even if a player in a set $T_k \in T$ has more preferable houses in a higher set T_h , those houses are exchanged within T_h in a mutually beneficial way. In addition, from the definition of a PMSS, no player in T_h has an incentive to trade with a player in lower sets T_k with $k > h$. Since the strict core is a wdom stable set, segmentability is also a sufficient condition for the existence of a wdom stable set. Quint and Wako (2004) gave a polynomial-time algorithm to examine segmentability of a house barter market. We show an example of a house barter market with segmentability.

Example 9 Let $M_3 = (N, R)$ be the market with the player set $N = \{1, 2, 3, 4, 5, 6\}$ and the following preference profile:

- 1) $2P_13P_15P_14P_11P_16$,
- 2) $1I_23P_24P_26P_25P_22$,
- 3) $1P_32P_33P_34P_35P_36$,
- 4) $2P_45P_46P_43P_44P_41$,
- 5) $1I_54P_55P_53P_56P_52$,
- 6) $3P_66P_61P_62P_64P_65$.

Although M_3 has two PMSSs, $T = \{T_1 = \{1, 2, 3\}, T_2 = \{4, 5\}, T_3 = \{6\}\}$ and $T' = \{T'_1 = \{1, 2, 3\}, T'_2 = \{6\}, T'_3 = \{4, 5\}\}$, the differences are only in the orders of sets in T and T' . In this market, the set $K = \{(2, 3, 1, 5, 4, 6)\}$ is the strict core.

The house barter market was also discussed by Moulin (1995) from a wide perspective of cooperative microeconomics and game theory. Recent studies have shown that interesting economic implications are derived from this market model when we assume the farsighted von Neumann-Morgenstern stability.

Farsighted Stable Sets in a General Setting

The Model

We have so far looked at stable sets using a domination relation that involved only a one-step deviation by a coalition. Harsanyi (1974) argued that these stable sets use domination relations that do not take into account possible subsequent deviations by some other coalitions and defined a new domination relation called indirect domination. Chwe (1994), based on this critique, modified the indirect domination concept based on a version laid out in the postscript of Harsanyi (1974), and laid out a general framework that includes both noncooperative and cooperative game models on which these concepts can be defined. Greenberg (1990) had already laid down the groundwork to apply von Neumann-Morgenstern stable sets to models outside of the characteristic function form games, such as strategic form games. (The theory of social situations, as laid out in Greenberg (1990), starts with an abstract framework, which Greenberg calls a situation, where the rules of the situation have been defined.

Throughout the many different situations that it covers, including strategic form games, the solution concept is kept constant, namely, that of von Neumann-Morgenstern stability.) Chwe's model is a simplification of one of the many models – or as Greenberg called them situations – introduced there. Farsighted stable sets are defined as von Neumann-Morgenstern stable sets defined by using Chwe's indirect domination relation. To distinguish between farsighted stable sets and the stable sets discussed earlier, we will refer at times to the latter as classical stable sets.

While the focus of this section is on farsighted stable sets, it should be noted that Harsanyi's original version of indirect domination has also been used in the literature.

Greenberg et al. (2002) used Harsanyi's indirect domination to define a stable set, which they call a sophisticated stable set, and applied them to an exchange economy. They show that there is a one-to-one correspondence property between the sophisticated stable sets defined on the set of payoffs in the economy and the sophisticated stable sets defined on the set of allocations. This property is shared by the core but not the classical stable set.

Chwe defines a game as the collection of the following primitives: $(N, X, (\preceq_i)_{i \in N}, (\rightarrow_S)_{S \subseteq N, S \neq \emptyset})$ where $N = \{1, 2, \dots, n\}$ represents the set of players, X is the set of possible outcomes, and $\preceq_i$ is player i 's preferences over the set X . These three elements are related to the components of a strategic form game, while the fourth component may not be as familiar. For each nonempty subset of players, $S \subseteq N$, $\rightarrow_S$ is a binary relation on X called the enforceability relation, or the effectiveness relation, where $x \rightarrow_S y$ indicates that when the status quo is x , the coalition S can induce outcome y by themselves. That is, the enforceability relation contains information about what coalitions can do. As we will see in the subsequent sections, this model is general enough that games in the game theoretic literature can be formulated in this manner. It will also be apparent that how the enforceability condition is defined affects the definition of indirect domination and the farsighted stable set.

Let x and y be two outcomes in X . We say that x indirectly dominates y and denote this by $x \gg y$

if there exists a sequence of outcomes $y = x^0, x^1, \dots, x^p = x$ and coalitions $S^1, S^2, \dots, S^p$ such that for each $j = 1, 2, \dots, p$, (i) $x^{j-1} \rightarrow_{S^j} x^j$ and (ii) $x^{j-1} \prec_{S^j} x^j$. Condition (i) states that each move from x^{j-1} to x^j is a feasible one by the coalition S^j , and these moves start from y and end at x . Condition (ii) states that all these coalitions are better off at the final outcome than at the outcome when they make their move. Note that it needs not be the case that all members in, for example, S^j are better off in outcome x^j than in x^{j-1} .

The notation $\prec_S$ denotes the preference relation of the coalition S and is implicitly assumed that it can be represented in terms of the preferences of the individual, $\prec_i$. In practice, $x \prec_S y$ if $x \prec_i y$ for all $i \in S$, or $x \prec_S y$ if $x \prec_i y$ for all $i \in S$ and $x \prec_i y$ for some $i \in S$. Unless specified otherwise, in the parts that follow, we employ the former interpretation for the notation $\prec_S$ and call the indirect domination using this version as the usual indirect domination. When we use the usual indirect domination in a given circumstance, we will not write out the definition of indirect domination. For those relations that are not usual, the indirect domination relation will be defined explicitly.

We hereupon remark that in the definition of indirect domination, it is implicitly assumed that joint moves by groups of players are neither once-and-for-all nor binding, i.e., some players in a deviating group may later make another move with players within or even outside the group. In some models, this assumption may not be valid, and in those circumstances, we impose restrictions on the relation $\rightarrow_S$. This is done in noncooperative strategic form games, in which coalitional deviations are not allowed, and in network formation games and matching games, in which deviations are conducted through only pairs or singletons.

When $p = 1$ in the definition of indirect domination, we simply say that x directly dominates y , which is denoted by $x \succ^d y$. When we want to specify a deviating coalition, we say that x directly dominates y via coalition S , which is denoted by $x \succ_S^d y$. This form of direct domination originates from the theory of social situations in Greenberg (1990) and allows us to define the myopic domination relation to strategic form games and other

game forms outside of the games in characteristic function form. We call the stable set defined by the relation $\succ^d$ as the myopic stable set. The myopic stable set is very similar to the classical stable set, except that the former is built off a model with an explicitly defined enforceability relation, while the latter is built off a domination relation that is implicitly myopic. See Kawasaki (2010) for details in the house barter game and Herings et al. (2017) for two-sided matching where those two concepts can be different. We mention briefly some other domination relations closely related to indirect domination. In Harsanyi's original definition of indirect domination, the first condition (i) is replaced by the following condition: (i') for each j , x^{j-1} is directly dominated by x^j . Therefore, Harsanyi's indirect domination is stronger than the indirect domination relation defined in Chwe (1994). Another domination relation can also be defined by using just condition (i') and without condition (ii). Page and Wooders (2009) used the term path dominance to describe this domination relation in the network formation model.

The Largest Consistent Set and the Largest Farsighted Conservative Stable Set

Because this entry focuses entirely on the stable set of von Neumann and Morgenstern, the majority of what follows focuses on stable sets defined on indirect domination relations defined in several models. However, Chwe (1994) also defined a solution concept called the largest consistent set. To obtain a better perspective on this solution concept, we first introduce an equivalent formulation of a farsighted stable set.

Let $K \subseteq X$ be a set of outcomes and for each $x \in X$, define $K_x = \{y \in K \mid y = x \text{ or } y \gg x\}$ to be the set of "likely outcomes" in K when the status quo is x . Then, K is a farsighted stable set if and only if the following two conditions are satisfied (see Diamantoudi and Xue (2003) and Xue (1998) for details):

- $x \in K \Rightarrow$ There do not exist $y \in X$ and $S \subseteq N$ with $x \rightarrow_S y$ such that for some $z \in K_y$, $z \succ_S x$.
- $x \notin K \Rightarrow$ There exist $y \in X$ and $S \subseteq N$ with $x \rightarrow_S y$ such that for some $z \in K_y$, $z \succ_S x$.

This characterization sheds light to the optimistic behavior that is implicitly assumed in the farsighted stable set. The deviating coalition S carries out the deviation if there is one outcome that is a likely outcome that makes the coalition S better off.

On the other hand, the consistent set defined in Chwe (1994) assumes a more conservative behavior in the following sense. When referring to the set K_y in the above definition, the phrase “for some” is replaced by “for all.”

Formally, $K \subseteq X$ is a consistent set if the following conditions hold:

- $x \in K \Rightarrow$ There do not exist $y \in X$ and $S \subseteq N$ with $x \rightarrow_S y$ such that for all $z \in K_y$, $z \succ_S x$.
- $x \notin K \Rightarrow$ There exist $y \in X$ and $S \subseteq N$ with $x \rightarrow_S y$ such that for all $z \in K_y$, $z \succ_S x$.

The set $L \subseteq X$ is called the largest consistent set (LCS) if it itself is a consistent set and for every consistent set K , $K \subseteq L$. Chwe (1994) showed that this concept is well defined, that is, there exists one and only one such set of outcomes that can be called the LCS. Moreover, it was shown that the LCS contains all farsighted stable sets.

A careful look at the definition reveals the possibility that the empty set can be a consistent set for any environment, as the condition involving “for all K_y ” can be satisfied vacuously. This possibility can also lead to the possibility of the LCS being empty as well. Chwe (1994) provided a sufficient condition that was later weakened by Xue (1997), guaranteeing the nonemptiness of the LCS.

Another approach around the possible emptiness of the solution concept is to modify the phrase to “for all $z \in K_y$ such that $K_y \neq \emptyset$ ” in the two conditions. Such is the approach taken in Greenberg (1990) for defining a conservative version of the (myopic) stable set and in Diamantoudi and Xue (2003) for the conservative version of the farsighted stable set in coalition formation games, which they call a farsighted conservative stable set. They also define the analogue to the LCS called the largest farsighted conservative stable set (LFCSS). When such a set exists, it coincides with the LCS, because the only part that separates the two concepts is nonemptiness.

In practice, the LCS is generally too inclusive to give a meaningful prediction to the model. In almost all strategic form games considered in the following, the LCS is the set of all individually rational outcomes. One exception is the game in Kawasaki and Muto (2009), but that is because every outcome is individually rational. Therefore, the general focus on the sections that follow is on farsighted stable sets.

Applications of Farsighted Stable Sets in Strategic Form Games

Let $G = (N, (X_i)_{i \in N}, (u_i)_{i \in N})$ be a game in strategic form, where N is the set of players, X_i is the set of strategies for player i , and u_i is player i 's payoff function. In order to apply farsighted stable sets to this framework, we need to reformulate this strategic form game into the language of Chwe's framework.

The first three components are relatively straightforward; N is the same for both models, $X = X_1 \times X_2 \times \dots \times X_n$, and the ordinal preferences $\preceq_i$ can be obtained from the payoff functions very easily. The enforceability relation is given by the following: for two outcomes x and y ,

$$x \rightarrow_S y \Leftrightarrow x_i = y_i \forall i \in N \setminus S$$

The enforceability relation for strategic form games states that when a coalition S deviates from the status quo x , the other players are assumed to be choosing their strategies in x . This assumption is closely related to the assumptions made in equilibrium concepts. Also, note that in this model, we allow coalitions to form freely. However, if we want to stick to the assumptions made in the concept of Nash equilibrium, we can similarly define an enforceability relation for such situations by the following:

$$x \rightarrow_S y \Leftrightarrow x_i = y_i \forall i \in N \setminus S \text{ and } |S| = 1$$

An indirect domination relation based on this enforceability relation can be defined. A stable set defined by this indirect domination is called a noncooperative farsighted stable set. Compared to myopic stable sets, in most games, farsighted stable sets give much sharper insights into


players' behavior in economic, political, and social situations. In the following, we first review the results for prisoner's dilemma games. For these games, the main result is that in most cases, only Pareto-efficient outcomes can be supported by farsighted stable sets. Next, we look at a public good provision game where the production level of the public good is either 0 or 1, and players either choose whether to contribute or not. This game is closely related to the prisoner's dilemma, but the results differ greatly as almost all strictly individually rational outcomes can be supported by a farsighted stable set. Then, we review the results for duopoly markets, which are closely related to the results in prisoner's dilemma.

Prisoner's Dilemma

To make the discussion as clear as possible, we will focus on a particular example of the prisoner's dilemma, which is given below. Similar results hold in general prisoner's dilemma games.

Prisoner's Dilemma:

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	4,4	0,5
	Defect	5,0	1,1



We first present a farsighted stable set derived when two players use only pure strategies. For shorthand, let C denote Cooperate and D denote Defect. In this case, the set of strategy combinations is $X = \{(C, C), (C, D), (D, C), (D, D)\}$, where in each combination, the former (resp. the latter) is player 1's (resp. 2's) strategy.

A myopic stable set does not exist in this game. On the other hand, the singleton $\{(C, C)\}$ is the unique farsighted stable set with respect to $\gg$. To see this, note that $(C, C) \gg (C, D), (C, C) \gg (D, C)$

via the sequence $(C, D) \rightarrow_1 (D, D) \rightarrow_{1,2} (C, C)$ and $(D, C) \rightarrow_2 (D, D) \rightarrow_{1,2} (C, C)$, respectively. Moreover, no other indirect domination relation exists. Hence, if the two players are farsighted and make a joint but not binding move, the farsighted stable set succeeds in showing that cooperation of the players results in the unique stable outcome.

We now study the farsighted stable sets in the mixed extension of the prisoner's dilemma, i.e., the prisoner's dilemma with mixed strategies played. Let $X_1 = X_2 = [0, 1]$ be the sets of mixed strategies of players 1 and 2, respectively, and let $t_1 \in X_1$ (resp. $t_2 \in X_2$) denote the probability that player 1 (resp. 2) plays Cooperate. It is easily seen that the minimax payoffs to players 1 and 2 are both 1 in this game. We say that a strategy combination is individually rational (resp. strictly individually rational) if both players' payoffs are at least (resp. exceed) their minimax payoffs. We then have the following theorem in Suzuki and Muto (2000).

Theorem 14 Let

$$T = \{(t_1, t_2) \mid 1/4 < t_1 \leq 1, t_2 = 1\} \cup \{(t_1, t_2) \mid t_1 = 1, 1/4 < t_2 \leq 1\},$$

and define the singleton set $K^1(t_1, t_2) = \{(t_1, t_2)\}$ for each $(t_1, t_2) \in T$. Let $K^2 = \{(0, 0), (1, 1/4)\}$ and $K^3 = \{(0, 0), (1/4, 1)\}$. Then the sets K^2 and K^3 and the singleton sets $K^1(t_1, t_2)$ with $(t_1, t_2) \in T$ are the farsighted stable sets of the mixed extension of the prisoner's dilemma, and there are no other types of farsighted stable sets. (The cutoff probabilities in the set T pertain to the example above, and these values may depend on the payoffs of the prisoner's dilemma. The overall characterization result of the farsighted stable sets nonetheless holds in general prisoner's dilemma games.)

This theorem shows that if the two players are farsighted and make a joint but not binding move in the prisoner's dilemma, then essentially a single Pareto-efficient and strictly individually rational strategy combination results as a stable outcome, i.e., $K^1(t_1, t_2)$. We, however, have two exceptional cases as shown by the sets K^2 and K^3 in which (D, D) could be stable together with one Pareto-

efficient point at which one player gains the same payoff as in (D, D) .

n-Person Prisoner's Dilemma

We consider an n-person prisoner's dilemma. Let $N = \{1, \dots, n\}$ be the player set. Each player i has two strategies: C (Cooperate) and D (Defect). Let $X_i = \{C, D\}$ for each $i \in N$. Sometimes, we will refer to a strategy combination as a state. For each coalition $S \subset N$, let $X_S = \prod_{i \in S} X_i$ and $X_{-S} = \prod_{i \in N \setminus S} X_i$. Let x_S and x_{-S} denote generic elements of X_S and X_{-S} , respectively. Player i 's payoff depends not only on his/her strategy but also on the number of other cooperators. Player i 's payoff function $u_i : X \rightarrow \mathcal{R}$ is given by $u_i(x) = f_i(x_i, h)$, where $x \in X$, $x_i \in X_i$ (player i 's choice in x), and h is the number of players other than i playing C . We call the strategic form game thus defined an n-person prisoner's dilemma game.

To make the arguments simple, we assume that all players are homogeneous and each player has an identical payoff function. That is, f_i 's are identical and simply written as f unless any confusion arises. We assume the following properties on the function f .

Assumption 1

1. $f(D, h) > f(C, h)$ for all $h = 0, 1, \dots, n-1$.
2. $f(C, n-1) > f(D, 0)$.
3. $f(C, h)$ and $f(D, h)$ are strictly increasing in h .

Property (1) states that every player prefers playing D to playing C regardless of which strategies other players play. Property (2) states that if all players play C , then each of them gains a payoff higher than the one in $(D, \dots, D)$. Property (3) states that if the number of cooperators increases, every player becomes better off regardless of which strategy is played.

It holds from Property (1) that $(D, \dots, D)$ is the unique Nash equilibrium of the game. Here for $x, y \in X$, we say that y is Pareto superior to x if $u_i(y) \geq u_i(x)$ for all $i \in N$ and $u_i(y) > u_i(x)$ for some $i \in N$. The state $x \in X$ is said to be Pareto efficient if there is no $y \in X$ that is Pareto superior to x . By Property (2), $(C, \dots, C)$ is Pareto superior

to $(D, \dots, D)$. Together with Property (3), $(C, \dots, C)$ is Pareto efficient.

Given a state x , we say that x is individually rational if for all $i \in N$, $u_i(x) \geq \min_{y_{-i} \in X_{-i}} \max_{y_i \in X_i} u_i(y)$. If a strict inequality holds, we say that x is strictly individually rational. From (1), (3) of Assumption 1 $\min_{y_{-i} \in X_{-i}} \max_{y_i \in X_i} u_i(y) = f(D, 0)$.

The following theorem from Suzuki and Muto (2005) shows that any strategy combination that is strictly individually rational and Pareto efficient is itself a singleton farsighted stable set. Moreover, there are no other farsighted stable sets except in the rarity that there exists a strategy combination that is Pareto efficient and individually rational but not strictly individually rational. In such a case, we have one more farsighted stable set, which also includes the outcome $(D, \dots, D)$. To state the result, we define the set $C(x) = \{i \in N \mid x_i = C\}$ as the set of players who choose the strategy C in $x \in X$.

Theorem 15 For the n -person prisoner's dilemma game, if x is a strictly individually rational and Pareto-efficient state, then $\{x\}$ is a farsighted stable set. Moreover, there are no other types of farsighted stable sets except in the following situation. If there exists a number s^* such that $f(C, s^* - 1) = f(D, 0)$ and each strategy combination y with $|C(y)| = s^*$ is Pareto efficient, then there exists exactly one more farsighted stable set given by $\{x' \in X \mid |C(x')| = s^*\} \cup \{(D, \dots, D)\}$.

Provision of Discrete Public Goods

A common economic application of the prisoner's dilemma is the provision of public goods. Consider a simple game in which every player has only two strategies: to "contribute (C)" and to "not contribute (D).". Contributing to the production of a public good comes with a cost. Suppose that a positive amount of the public good can be produced even if only one player contributes. Because everyone can enjoy the benefits of the public good once it is produced, every player has a dominant strategy to not contribute. Thus, this situation would be modeled as a prisoner's dilemma game.

Now, suppose that the public good is provided in discrete amounts and that it requires a minimum of r^* players contributing to produce the good where $r^* \geq 2$. Using the function f defined in the previous subsection, we consider a game that satisfies the following conditions. Recall that h represents the number of other players choosing C .

Assumption 2

1. $f(D, h) > f(C, h)$ for all $h \geq r^*$, $f(D, h) = f(C, h)$ for all $h \leq r^* - 2$, $f(D, r^* - 1) < f(C, r^* - 1)$
2. $f(C, n - 1) > f(D, 0)$
3. $f(C, 0) = \dots = f(C, r^* - 2) < f(C, r^* - 1) = \dots = f(C, n - 1)$ and $f(D, 0) = \dots = f(D, r^* - 1) < f(D, r^*) = \dots = f(D, n - 1)$

The first condition states that choosing D is strictly better than C if the number of other contributors is at least r^* , which is when the public good can be produced without any additional contribution. Therefore, there is some incentive to free ride, but choosing D is strictly worse than C when there are exactly $r^* - 1$ other players choosing C . We also assume that choosing C or D leads to identical results when the public good is not produced. The second condition is the same as the prisoner's dilemma. The third condition states that the benefit of choosing C or D depends on the number of other contributors only through whether the public good is produced or not. Thus, instead of strictly increasing in h , the function f is mostly flat with respect to h and takes a jump at the threshold amount at which the public good is produced or not.

Just as in the prisoner's dilemma, every strategy combination that is Pareto efficient and strictly individually rational constitutes a singleton farsighted stable set. Unlike the prisoner's dilemma, however, many other types of farsighted stable sets may exist. In fact, every outcome that is strictly individually rational but not Pareto efficient, except for $(C, C, \dots, C)$, is included in some farsighted stable set. See Kawasaki and Muto (2009) for details.

Theorem 16 Let k be a positive integer such that $r^* + k \neq n$, and let $x \in X$ be such that the number of players choosing C in x is $r^* + k$. Then, there exists a farsighted stable set that includes x .

Duopoly Market Games

We consider two types of duopoly markets: Cournot quantity-setting duopoly and Bertrand price-setting duopoly. For simplicity, we will consider a simple duopoly model in which firms' cost functions and the market demand function are linear. Similar results, however, hold in more general duopoly models.

There are two firms, 1 and 2, each producing a homogeneous good with the same marginal cost $c > 0$. No fixed cost is assumed.

1. **Cournot duopoly:** Firms' strategic variables are their production levels. Let x_1 and x_2 be production levels of firms 1 and 2, respectively. The market price $p(x_1, x_2)$ for x_1 and x_2 is given by

$$p(x_1, x_2) = \max(a - (x_1 + x_2), 0),$$

where $a > c$. We restrict the domain of production of both firms to $0 \leq x_i \leq a - c, i = 1, 2$. This is reasonable since a firm would not overproduce to make a nonpositive profit. When x_1 and x_2 are produced, firm i 's profit is given by

$$\pi_i(x_1, x_2) = (p(x_1, x_2) - c)x_i.$$

Thus, Cournot duopoly is formulated as the following strategic form game

$$G^C = (N, \{X_i\}_{i=1,2}, \{\pi_i\}_{i=1,2}),$$

where the player set is $N = \{1, 2\}$; each player's strategy set is a closed interval between 0 and $a - c$, i.e., $X_1 = X_2 = [0, a - c]$; and their payoff functions are $\pi_i, i = 1, 2$. Let $X = X_1 \times X_2$. The joint profit of two firms is maximized when $x_1 + x_2 = (a - c)/2$.

2. **Bertrand duopoly:** Firms' strategic variables are their price levels. Let

$$D(p) = \max(a - p, 0)$$

be the market demand function. Then the total profit at p is

$$\prod(p) = (p - c)D(p).$$

We restrict the domain of price level p of both firms to $c \leq p \leq a$. This assumption is also reasonable since a firm would avoid a negative profit. The total profit $\prod(p)$ is maximized at $p = (a + c)/2$, which is called the monopoly price. Let p_1 and p_2 be prices chosen by firms 1 and 2, respectively. We assume that if firms' prices are equal, then they share equally the total profit; otherwise, all sales go to the lower-pricing firm of the two. Thus, firm i 's profit is given by

$$\rho_i(p_i, p_j) = \begin{cases} \prod(p) & \text{if } p_i < p_j \\ \prod(p)/2 & \text{if } p_i = p_j \\ 0 & \text{if } p_i > p_j \end{cases}$$

for $i, j = 1, 2, i \neq j$.

Hence, Bertrand duopoly is formulated as the strategic form game

$$G^B = (N, \{Y_i\}_{i=1,2}, \{\rho_i\}_{i=1,2}),$$

where $N = \{1, 2\}$, $Y_1 = Y_2 = [c, a]$ and $\rho_i (i = 1, 2)$ is i 's payoff function. Let $Y = Y_1 \times Y_2$.

It is well known that a Nash equilibrium is uniquely determined in either market: $x_1 = x_2 = (a - c)/3$ in the Cournot market and $p_1 = p_2 = c$ in the Bertrand market. The following theorem holds for the farsighted stable sets in a Cournot duopoly.

Theorem 17 Let $(x_1, x_2) \in X$ be any strategy pair with $x_1 + x_2 = (a - c)/2$. Then the singleton $\{(x_1, x_2)\}$ is a farsighted stable set. Furthermore, every farsighted stable set is of the form $\{(x_1, x_2)\}$ with $x_1 + x_2 = (a - c)/2$ and $x_1, x_2 \geq 0$.

As mentioned before, any strategy pair (x_1, x_2) with $x_1 + x_2 = (a - c)/2$ and $x_1, x_2 \geq 0$ maximizes the firms' joint profit. This suggests that the von Neumann-Morgenstern stability together with firms' farsighted behavior yields joint profit maximization even if firms' collaboration is not binding.

As for Bertrand duopoly, we have the following theorem, which claims that the monopoly price pair is itself a farsighted stable set, and no other farsighted stable set exists. Therefore, the

von Neumann-Morgenstern stability together with firms' farsighted behavior attains efficiency (from the standpoint of firms) also in Bertrand duopoly. We refer the reader to Suzuki and Muto (2006) for the details.

Theorem 18 Let $p = (p_1, p_2)$ be the pair of monopoly prices, i.e., $p_1 = p_2 = (a + c)/2$. Then the singleton $\{p\}$ is the unique farsighted stable set.

Some General Results for Strategic Form Games

We have shown in the previous parts some applications of farsighted stable sets to the prisoner's dilemma and duopoly games. Below we explain very briefly some general findings regarding farsighted stable sets of strategic form games.

Kawasaki (2015) considers general strategic form games with two players and provides a sufficient condition for a strictly individually rational and Pareto-efficient outcome to be a singleton farsighted stable set. The term individually rational used in this article is defined using the minimax value instead of the maximin value (obtained by interchanging the max and min operations), but in the games introduced here, the maximin value and the minimax value coincide, which may not be the case in general. We have seen that similar results hold for the prisoner's dilemma game with possibly more than two players, but Kawasaki (2015) notes that similar results do not generally hold in those instances.

Hirai (2017) shows that these results can be recovered for three or more players under a certain class of games, connected with the concept of punishment dominance in Nakayama (1998). Informally, a strategy is punishment dominant toward the set of other players if that particular strategy brings the payoffs of the other players collectively lower than any other strategy. In the prisoner's dilemma and in the public goods provision model, D is a punishment dominant strategy. It is also known in that in these classes of games, the maximin and the minimax values coincide.

Further Research on (Myopic and Farsighted) Stable Sets in Strategic Form Games

The literature on farsighted stable sets introduced here is by no means exhaustive. Here we briefly list other papers that studied farsighted stable sets

in strategic form games and mention a few papers that studied myopic stable sets in strategic form games. Masuda (2002) analyzed farsighted stable sets in average return games and obtain similar results to the prisoner's dilemma. Diamantoudi (2005) defined a farsighted stable set for a cartel price leadership model and proves its existence nonconstructively; on the other hand, in order to characterize these farsighted stable sets, Kamiyo and Muto (2010) reformulated the model into a strategic form game similar to the prisoner's dilemma and obtained similar results to Suzuki and Muto (2005). Their results imply that some of the conditions in the payoff function can be weakened to obtain the results of Suzuki and Muto (2005). In location games, Shino and Kawasaki (2012) showed the existence of a farsighted stable set that supports two different location profiles: minimum differentiation and local monopoly. Kawasaki et al. (2015) apply the solution concept to an international trade model similar to that of Oladi (2005) and Nakanishi (1999) with two countries where each chooses a tariff rate on imports. The game resembles closely to that of the duopoly games introduced above, and they obtain very similar results in that model.

The enforceability relation for strategic form games states that when a coalition S deviates from the status quo x , the other players are assumed to be choosing their strategies in x . This assumption is very closely related to the assumptions made in equilibrium concepts. Also, note that in this model, we allow coalitions to form freely. However, if we want to stick to the assumptions made in noncooperative game theory, we can similarly define an enforceability relation that explicitly forbid joint moves by two or more players.

An indirect domination relation based on the enforceability relation with the restriction explained in the previous paragraph can be defined. A stable set defined by this indirect domination is called a noncooperative farsighted stable set. Nakanishi (2009) considered the noncooperative farsighted stable sets of the prisoner's dilemma and shows the existence of a unique noncooperative farsighted stable set, which includes the outcome in which all players choose D and some Pareto-efficient outcome.

Myopic stable sets for the prisoner's dilemma and the duopoly markets have been applied in Nakanishi (2001) and in Muto and Okada (1996, 1998), respectively. Myopic stable sets have also been applied to other economic models such as international trade. Nakanishi (1999) analyzed export quota games, while Oladi (2005) considered tariff retaliation games. For further studies on other solution concepts related to stable sets for strategic form games, see Kaneko (1987) and Mariotti (1997).

Farsighted Stable Sets in Cooperative Games

In this section, we review some results on farsighted stable sets and related solution concepts to games other than strategic form games. Such games include not only characteristic function form games, both transferable utility (TU) and nontransferable utility (NTU), but also network formation games, coalition form games, and matching markets. The literature on these models has grown very rapidly.

Part of the challenge in applying Chwe's framework in this setting is defining an appropriate enforceability condition. Because unlike strategic form games, some cooperative game models do not specify every detail about what players can or cannot do. For example, when a coalition deviates, nothing is stated about the outcome pertaining to the other players. This information was not necessary in defining the domination relation for the core and classical stable set. In the following, we look at several models from the perspective of how the enforceability relations are defined.

Characteristic Function Form Games and Coalitional Sovereignty

This section focuses on the farsighted stable sets of characteristic function form games. Recall that the original stable sets were first defined for characteristic function form games. The results introduced here then allow us to make some comparisons between the classical stable set and the farsighted stable set. First, we introduce some results from Beal et al. (2008), which analyzed farsighted stable sets of transferable utility (TU) characteristic function form games. Then,

we review a result from Bhattacharya and Brosi (2011), which focused on the nontransferable utility (NTU) characteristic function form games:

1. TU games: Let (N, v) be a TU game. We first formulate this in terms of Chwe's framework. The set of outcomes X corresponds to the set of imputations, typically labeled by $I(N, v)$. The preferences of a player i over the set of outcomes is such that player i likes the imputation that gives this player a higher amount. That is, for two imputations x, y , $x \prec_i y$ if and only if $x_i < y_i$. As for the enforceability condition, the relation $\rightarrow_S$ is defined in the following way:

$$x \rightarrow_S y \Leftrightarrow \sum_{i \in S} y_i \leq v(S)$$

Beal et al. (2008) show the existence of farsighted stable sets in TU game and in fact characterize the farsighted stable sets of TU games (more details to be laid out later).

2. NTU games: Bhattacharya and Brosi (2011) also established the existence of a farsighted stable set for NTU games under mild conditions. The relation $\rightarrow_S$ is defined in a way similar to that in TU games. Formally, let (N, V) be an NTU game, where the characteristic function V is now a set-valued function that maps to each coalition S a region of the n -dimensional payoff space representing the feasible payoffs of a coalition S .

To formulate NTU games in terms of the primitives of Chwe's framework, let $X = V(N)$; preferences are defined in the same way as TU games, and the relation $\rightarrow_S$ is defined in the following way:

$$x \rightarrow_S y \Leftrightarrow y \in V(S)$$

Let b_i denote the maximum payoff for i that can be attained in the set $V(\{i\})$ and define a vector $x \in V(N)$ to be individually rational if $x_i \geq b_i$ for all $i \in N$. Bhattacharya and Brosi (2011) showed that there exists a farsighted stable set if the set of individually rational vectors is bounded, thereby

generalizing the existence portion of the result by Beal et al. (2008).

In both TU games and NTU games, the respective papers derive what can be seen as existence results suggesting that farsighted stable sets are more likely to exist than classical stable sets. The main reason for this is that it is easier for one outcome to indirectly dominate the other. One possibly problematic byproduct of such indirect domination relation is the seeming arbitrariness of the imputations supported by farsighted stable sets. In particular, for TU games, Beal et al. (2008) give the following characterization result.

Theorem 19 Suppose that for a TU game (N, v) , $v(N) > \sum_{i \in N} v(\{i\})$ holds. Then, all farsighted stable sets are singleton sets containing an imputation x such that for some coalition S , $\sum_{i \in S} x_i \leq v(S)$ and $x_i > v(\{i\})$ for all $i \in S$.

This result then implies the surprising result that for superadditive games in which the Shapley value is not in the core, the Shapley value constitutes a singleton farsighted stable set. (This result does not exclude the possibility that when the Shapley value is in the core, it is a farsighted stable set. We refer the reader to Beal et al. (2008) for details.) Another implication is that imputations in the interior of the core cannot be supported by a farsighted stable set.

Ray and Vohra (2015a) argue that the peculiarity of these farsighted stable sets stems from the fact that the enforceability condition does not satisfy what they call *coalitional sovereignty*. Specifically, taking the TU game as an example, when $x \rightarrow_S y$ holds for some coalition S , the only condition that needs to be satisfied is $\sum_{i \in S} y_i \leq v(S)$. Therefore, y_j for $j \notin S$ can be chosen arbitrarily by players in S for their liking. In fact, this situation is unavoidable, since imputations must satisfy the condition $\sum_{i \in N} x_i = v(N)$, so that if members in S were to move to an imputation that benefits each member, then someone outside S must receive a smaller amount.

In the words of Ray and Vohra (2015a), *coalitional sovereignty* is violated in such models, and the main purpose of their paper is to construct a model that respects coalitional sovereignty that

is connected to games in characteristic function form. In the following, we briefly summarize their model. First, an outcome is defined by a pair $x = (u, \pi)$ where $u \in \bar{V}(N)$ represents the set of payoff vectors in $V(N)$ that are Pareto efficient and π is a partition of N that represents the coalition structure. For each state x , its payoff vector and coalition structure may be denoted by $u(x) = (u_i(x))_{i \in N}$ and $\pi(x)$, respectively. The specifics of a coalition formation model will be explained in the later sections, but here such formality is not needed in explaining the model.

Instead of defining an explicit enforceability condition for this framework, Ray and Vohra (2015a) impose the following conditions that should be satisfied for an enforceability relation:

1. If $x \rightarrow_T y$ and $S \in \pi(x)$ with $T \cap S = \emptyset$, then $u_i(x) = u_i(y)$ for all $i \in S$.
2. For every $x \in X$, $T \subseteq N$, and $v \in \bar{V}(T)$, there is a $y \in X$ such that $x \rightarrow_T y$, $T \in \pi(y)$, and $u_i(y) = v_i$ for all $i \in T$.

The set $\bar{V}(T)$ is defined similarly as $\bar{V}(N)$ as the set of feasible vectors that are Pareto efficient for T .

Direct domination can also be defined in this framework, and from this relation, the core can be defined by the set of outcomes x that are not directly dominated. If $x = (u, \pi)$ is in the core, then u is said to be in the coalition structure core. While farsighted stable sets of the original characteristic function game, especially the TU game, can be quite different from the core, Ray and Vohra (2015a) instead show that in their framework, an allocation vector is very closely related to the coalition structure core. The version of the results we present here is from Ray and Vohra (2015b). The original version in Ray and Vohra (2015a) gives a single condition that is both necessary and sufficient conditions called separability for a payoff vector to be a singleton farsighted stable set. Note that their result does not exclude the possibility of non-singleton farsighted stable sets.

Theorem 20 In the model of Ray and Vohra (2015a), the following statements hold:

1. If u is in the interior of the coalition structure core, then u is a singleton farsighted stable set. (To be precise, since the farsighted stable set is defined on the set of (u, π) , there should be an underlying coalition structure attached to u . The term singleton should then be interpreted as single payoffs as π can be different.)
2. If u is not in the coalition structure core, then u is not a singleton farsighted stable set.

Coalitional sovereignty is a condition that was not granted to coalitions in the classical TU model. In the following sections, we look at models in which coalitional sovereignty mostly is satisfied.

Network Formation Games

Here we focus on the network formation model formally described in Jackson and Wolinsky (1996). Let $N = \{1, 2, \dots, n\}$ be the set of players, which is also seen as the set of nodes in a graph. A network g on N is defined as the set of edges where $ij \in g$ is shorthand for the players $i, j \in N$ linked together by the edge (i, j) . The networks we introduce here do not distinguish between the edge from i to j and the edge from j to i and thus are called undirected networks. In some models, the direction of the links may matter, and these networks are called directed networks. However, we will only consider undirected networks in this section.

For $ij \notin g$, denote the network of adding the edge (i, j) to g by $g + ij$. Similarly, for $ij \in g$, denote the network of deleting the edge ij from g by $g - ij$. In the network formation literature, it is assumed that the edge ij can be deleted by either i or j without the other player's consent, while adding ij requires the consent of both players.

For simplicity, we assume that each player has preferences, denoted by the binary relation $\prec_i$, over the possible networks that can be formed. This setup contrasts to the original framework of Jackson and Wolinsky (1996). In their framework, they first defined a value function that represents the total worth of each network, just as the characteristic function for TU games represented the total worth for each coalition. Then, they defined an allocation function which defines how the

value would be split among the players depending upon which network is formed.

In terms of Chwe's framework, we now have the set of players, the set of outcomes being the set of possible networks, and the preferences that are defined directly on the networks. What remains to be defined is the enforceability relation $\rightarrow_S$. We introduce three versions of $\rightarrow_S$, as classified in Page and Wooders (2009), which are defined in the network formation literature. It should be noted that the original formulations in their paper are for directed networks, but they can be translated to undirected networks as well:

1. **Jackson-Wolinsky enforceability:** The main building blocks for the enforceability relation are singletons and pairs. Let $i, j \in N$ and g be a network. If $g' = g + ij$, then we have $g \rightarrow_{ij} g'$ where $i \neq j$. If $g' = g - ij$, we have $g \rightarrow_i g'$ and $g \rightarrow_j g'$, and for notational purposes, we allow $g \rightarrow_{ij} g'$. To form a link between two players, consent is needed from both players, while the deletion of a link can be carried out by just one player. Also, coalitional deviations of size greater than two are not allowed so that the only coalitions that can deviate are singletons and pairs. Therefore, $g \rightarrow_S g'$ cannot hold for any pair of networks g, g' when $|S| \geq 3$. This enforceability relation is taken from the concept of pairwise stability in Jackson and Wolinsky (1996).
2. **Jackson-van den Nouweland enforceability:** This enforceability relation, taken from the concept of strong stability defined by Jackson and van den Nouweland (2005), now allows coalitions of size greater than two to deviate so that multiple links can be formed and/or eliminated in one step. It is essentially the coalitional extension of the Jackson-Wolinsky enforceability relation. Formally, $g \rightarrow_S g'$ if and only if the following conditions are satisfied. If $ij \in g' \setminus g$, then $\{i, j\} \subseteq S$. If $ij \in g \setminus g'$, then $\{i, j\} \cap S \neq \emptyset$. The first condition states that the coalition S must contain all players involved in forming a link that was not there before. The second condition states that S must contain at least one of the players involved in the link that is destroyed.

3. **Bala-Goyal enforceability:** The enforceability condition is built off of the noncooperative model of network formation in Bala and Goyal (2000). The main feature of this enforceability relation is that links can be formed and severed by only one player so that no consent of the other player is needed to form a link. This relation seems a bit forceful for undirected networks, but for directed networks, because the link ij can be seen as the link from i to j , it may not be unnatural to think of models in which this link can be set up by i alone. However, it may be a bit of a stretch for i to have the power to form link ji , so this move is not allowed in the definition for directed networks. Moreover, coalitional deviations are not allowed.

An appropriate indirect domination relation can be defined for each enforceability relation. Page et al. (2005) defined a framework called supernetworks which translate these concepts visually into a directed network model defined in the following way. First, each node represents a network g , so that in essence, a network is built on top of a network, hence the terminology supernetwork. There are two types of directed edges, defined for each coalition S : a move edge and a preference edge. A move edge corresponding to a coalition S from network g to g' is drawn if $g \rightarrow_S g'$. A preference edge corresponding to S is drawn from g to g' if $g' \succ_S g$, depending upon how $\succ_S$ is defined.

The most commonly used of the three enforceability relations is the Jackson-Wolinsky enforceability, and for the remainder of the section, we employ this enforceability relation to define an indirect domination relation. Instead of defining indirect domination in the usual way in which $\succ_S$ is typically defined, we look at an indirect domination that uses the weaker notion of $\succ_{S-}$, i.e., $\succ_i$ for all $i \in S$ and $\succ_i$ for some $i \in S$ – as Jackson and Wolinsky (1996) defined their concept of pairwise stability using this weaker version. We give the formal definition below. Using the Jackson-Wolinsky enforceability relation, we say that a network g is indirectly dominated by g and is denoted by $g \ll g'$ if there exists a sequence of networks $g = g^0, g^1, \dots, g^p = g'$ and pairs

$(i_1, j_1), \dots, (i_p, j_p)$, where we include the possibility that for some k we have $i_k = j_k$; the following conditions are satisfied: (i) $g^{k-1} \rightarrow_{i_k, j_k} g^k$, (ii) $g^{k-1} \lesssim_l g'$ for all $l \in \{i_k, j_k\}$, and (iii) $g^{k-1} \prec_l g'$ for some $l \in \{i_k, j_k\}$.

In general, a stable set defined by the indirect domination described above may not exist. Herings et al. (2009) proposed an alternative farsighted solution concept, which we call the Herings-Mauleon-Vannetelbosch (HMV) farsighted stable set. To distinguish from the farsighted stable set defined in the earlier sections, we will at times refer to it as the vNM farsighted stable set. We present the formal definition in the following paragraphs, but we make one remark on the intuition behind the solution concept. The main modification in the HMV farsighted stable set is that internal stability is weakened to guarantee the existence of an HMV farsighted stable set. Due to the multiplicity of such sets, they then took the minimal set with respect to set inclusion that satisfies their modified internal stability and external stability. We once again use the notation $K_g = \{g' \in K \mid g' \gg g\} \cup \{g\}$ as the set of likely outcomes when the status quo is g . A set K is an HMV farsighted stable set if it satisfies the following conditions (where we allow $i=j$ in the following):

1. $g \in K \Rightarrow$ there does not exist $g' \notin K$ with $g \rightarrow_{i,j} g'$ and $g'' \in K_{g'}$ such that $g \lesssim_k g''$ holds for $k \in \{i, j\}$ and with strict preference for either i or j .
2. $g \notin K \Rightarrow$ there exists g' with $g \rightarrow_{i,j} g'$ and $g'' \in K_{g'}$ such that $g \lesssim_k g''$ holds for $k \in \{i, j\}$ and with strict preference for either i or j .
3. There is no set $K' \subsetneq K$ that satisfies 1 and 2.

The set of all networks satisfies the first two conditions vacuously. Therefore, because the set of networks is finite as long as the set of players is finite, there exists at least one HMV farsighted stable set. (This property has been one of the reasons why this version of the stable set has been used in the free trade agreement model of Zhang et al. 2013.)

Also, if the condition $g' \notin K$ is removed from the first condition, then the first two conditions are equivalent to the conditions specified by the vNM

farsighted stable set. Below are several facts that relate the two solution concepts that were proved in Herings et al. (2009).

Theorem 21 The following statements hold:

1. Every vNM farsighted stable set is an HMV farsighted stable set.
2. Every singleton HMV farsighted stable set is a vNM farsighted stable set.
3. If K is the unique HMV farsighted stable set, then it is also the unique vNM farsighted stable set.

In the various games introduced thus far, the implicit assumption was that players were sufficient farsighted in the sense that players can foresee the sequence of outcomes in an indirect domination of arbitrary (but finite) length. However, it may be the case that players that we observe may not have such unlimited foresight, and in the framework of network formation, a solution concept that takes into account limited foresight is called “horizon- k ” farsighted stable set defined in Herings et al. (2018). The motivation behind considering players with limited foresight originates from Kirchsteiger et al. (2016) which analyze by experiments whether players are indeed farsighted or myopic in a network formation game. They find that in many examples, players form a network that is supported by an intermediate level of farsighted behavior over myopic behavior. To our knowledge, this paper is the only one that looks at the issue of farsighted versus myopic through experiments, and this direction could be of interest going forward.

Coalition Formation Games

While in network formation games, the objective is to form bilateral links between two players, in coalition formation games, players partition themselves into coalitions so that players interact with other players multilaterally. The objective in this game is to divide players into a stable partition, where stability can be defined in many ways.

Let N be the set of players. A coalition structure is defined as simply a partition of N . That is, a coalition structure is given by $P = \{S_1, S_2, \dots, S_k\}$

where $S_i \cap S_j = \emptyset$ for all $i \neq j$ and $\bigcup_{i=1}^k S_i = N$. Let $\mathcal{P}$ denote the set of all coalition structures. Each player has a preference relation $\prec_i$ defined over $\mathcal{P}$. These two components constitute a game of coalition formation. The main objective is to look at coalition structures which satisfy some stability properties. In what follows, we first look at a special class of coalition formation games called hedonic games, in which each player only cares about which coalition he/she will be in and not about how the other players are partitioned. Then, we cover the results of games that are not hedonic and involve externalities.

Hedonic Games

Formally, hedonic games are coalition formation games in which for each player $i \in N$, the preference relation $\prec_i$ is defined over $P_i = \{S \subseteq N \mid i \in S\}$ instead of $\mathcal{P}$. The seminal papers on these games in this form are Banerjee et al. (2001) and Bogomolnaia and Jackson (2002), which defined several solution concepts, including the core. For a discussion of numerous concepts defined in this framework, see Sung and Dimitrov (2007). Barberá and Gerber (2003) defined a concept called durability that is defined recursively in a similar way that the coalition-proof Nash equilibrium (CPNE) of Bernheim et al. (1987) is defined. This concept does consider sequences of deviations but is not exactly a farsighted solution concept, as players in this model also care about the immediate consequences of their deviations.

Diamantoudi and Xue (2003) examined the farsighted conservative stable sets (defined earlier) in hedonic games and proved the existence of the largest conservative farsighted stable set (LFCSS). To define these concepts in this framework, we need to formulate the game in the language of Chwe's model. The set of players, the set of outcomes, and the preferences have already been defined. It remains to define the relation $\rightarrow_S$, for which there may be many ways of defining depending upon the model. We give the definition that was given in Diamantoudi and Xue (2003) below.

Let P and P' be two coalition structures and take a coalition $S \subseteq N$. Then, define the relation $\rightarrow_S$ by

$$P \rightarrow_S P' \Leftrightarrow P' = \{S\} \cup \{T \setminus S \mid T \in P, T \setminus S \neq \emptyset\}. \quad (1)$$

Although the formulation is simple, there are two important implicit assumptions. To make the point clearer, consider a subset S' that consists of those players who have left their coalitions. That is, if we let $P = \{S_1, S_2, \dots, S_k\}$, we can express S' as $S' = \bigcup_{i=1}^k T_k$, where for each j , $T_j \subseteq S_j$. Some of the T_k 's can be empty. Coalition S' can then join some remnant of the k coalitions in P to form a coalition S that would be part of the coalition structure P' . Because this deviation requires the consent of members in S and S only, this implies that no consent is needed when a group of players deviate, but consent is needed from the coalition which they wish to join.

Using the above relation $\rightarrow_S$, we can define the indirect domination and the direct domination relations. We also define the core as the set of coalition structures that are not directly dominated. The core is used as a benchmark solution concept when looking at the effects of introducing a farsighted solution concept.

There are other ways to define $\rightarrow_S$ based on the definition of the other stability concepts in the literature. We do not explain them in detail here as the main results of Diamantoudi and Xue (2003) hold for the enforceability relation in Eq. (1).

First, they show that in general, the LFCSS exists. This statement is a parallel result of the nonemptiness of the LCS in numerous environments.

Next, they focus on the domain of games in which the preferences of players are strict. Under this assumption, Diamantoudi and Xue (2003) proved the stability of a coalition structure in the (myopic) core. We state their result in the following theorem using slightly different terms from the original paper.

Theorem 22 Under the assumption of strict preferences, any coalition structure in the core constitutes a singleton farsighted (conservative) stable set.

Note that the word conservative is in parentheses, as any singleton set that is a farsighted stable set is by definition also a farsighted conservative

stable set. A corollary of the above fact is that the LFCSS contains the core.

A sufficient condition for the LFCSS to coincide with the core is that the hedonic game satisfies what is called the top coalition property, introduced in Banerjee et al. (2001). Informally, the top coalition property is such that there exists a coalition structure $\{S_1, S_2, \dots, S_k\}$ such that those in S_1 agree that S_1 is the most preferred coalition among subsets of N , those in S_2 agree that S_2 is the most preferred coalition among subsets of $N \setminus S_1$, etc. Under this condition, the coalition structure $\{S_1, S_2, \dots, S_k\}$ is the only coalition structure in the core. Diamantoudi and Xue (2003) showed that this coalition structure as a singleton is the LFCSS.

General Model

Diamantoudi and Xue (2007) considered farsighted solution concepts for general coalition formation games in which there are externalities. They first gave an alternative definition of the notion of equilibrium binding agreements (EBA), defined by Ray and Vohra (1997). The original definition of an EBA is defined recursively and involves a nestedness assumption on the coalitions that can deviate in the sequence of deviations, much like how the CPNE is defined for strategic form games.

Diamantoudi and Xue (2007) first formulated the EBA as an element of a stable set defined by a suitably defined domination relation, which they call R&V domination. This formulation allows the recursive nature of the EBA to be incorporated into the domination relation. Within this domination relation, the coalitional deviations that are considered are only those that involve a coalition splitting from one so that the resulting coalition structure now includes the deviating coalition and the remaining coalition. Thus, coalitions can only become smaller in this sequence of deviations, and in this sense, it is said that the deviations have a nested structure. Therefore, for the EBA to make sense, the starting coalition structure needs to be the grand coalition structure $\{N\}$.

By relaxing this nested structure when considering the EBA, Diamantoudi and Xue (2007) defined a new solution concept called the

extended EBA (EEBA). A coalition structure is an EEBA if it is an element of a stable set of an indirect domination relation very similar to what they had defined in their earlier paper (Diamantoudi and Xue 2003). This concept is built off of the EBA in that it considers sequences of deviations by coalitions, but unlike the EBA, the deviating coalitions no longer need to be nested and can also now merge in some steps of the deviation.

Diamantoudi and Xue (2007) then analyzed the relationship between the coalition structures that are in the EEBA and efficient coalition structures. In their result, they find a sufficient condition for the grand coalition to be an EEBA.

Herings et al. (2010), just as in Herings et al. (2009) for the network formation model, also considered an HMV farsighted stable set for the coalition formation games. Their results in this framework are very similar to those that they obtain for network formation games.

Funaki and Yamato (2014) took another approach and considered a different enforceability relation $P \rightarrow_S P'$. The relation defined earlier includes, in one step, disintegration of some coalitions when members of S leave the coalitions that they were part of in P and integration when forming the coalition $S \in P'$. Funaki and Yamato (2014) argued that these two moves should be treated separately, since in certain situations, disintegration and integration cannot both occur simultaneously as was assumed in the EEBA. Moreover, in their formulation, they only allow one disintegration or integration to occur in each step so that the number of coalitions either decreases by one by the integration of two separate coalitions or increases by one by the disintegration of one coalition into two separate coalitions. This way of interpreting $\rightarrow_S$ as representing one basic step is similar to the reasoning behind allowing only individual deviations in strategic form games. Formally, their enforceability relation can be stated as follows: $P \rightarrow_S P'$ if (i) $\{T \in P \mid T \subseteq N \setminus S\} = \{T' \in P' \mid T' \subseteq N \setminus S\}$; (ii) $S \in P'$ is such that $S = S_1 \cup S_2$ for some $S_1, S_2 \in P$ and $|P'| = |P| - 1$; and (iii) $S \in P'$ is such that there is some $T \in P$ with $T \setminus S \in P'$ and $|P'| = |P| + 1$.

The first condition simply states that the coalitions unaffected by the deviation by S are intact. The second condition states that the deviation by S must either result in two coalitions S_1 and S_2 merging to form S (the first case) or a coalition T splitting into two coalitions S and $T \setminus S$ (the second case). Hence, in the first case, the number of coalitions in P' must be one less than the number of coalitions in P , while in the second case, the number of coalitions in P must be one more than the number of coalitions in P .

Using the above relation, an indirect domination relation can be defined in the usual way. Funaki and Yamato (2014) focused on coalition structures with the property that each coalition structure indirectly dominates all other coalition structures. A coalition structure is said to be sequentially stable in that instance. Equivalently, each sequentially stable coalition structure constitutes a singleton stable set defined by the indirect domination relation that uses the stepwise enforceability relation defined in the previous paragraph.

They obtained a sufficient condition when the grand coalition as a coalition structure is sequentially stable and an algorithm to check that the grand coalition is a sequentially stable coalition structure.

Marriage Games and Roommate Games

In the next two sections, we revisit two models: the marriage game and the house barter game. Also in this section, we review results on roommate games, which can be seen as a generalized model of marriage games. The results for marriage games were first obtained by Mauleon et al. (2011), and the results for roommate games were shown by Klaus et al. (2011).

Let us first consider the marriage game (M, W, R) . All of the components in the Chwe's model, with the exception of $\rightarrow_S$, are apparent, so it remains to define the enforceability relation $\rightarrow_S$. Given two matchings μ and μ' and the coalition $S = \{i, j\}$, $\mu \rightarrow_{ij} \mu'$, if (i) $\mu'(i) = j$, (ii) $\mu'(k) = \mu(k)$ for $k \notin \{i, j, \mu(i), \mu(j)\}$, and (iii) $\mu'(k) = k$ for all other k .

Condition (i) states that i is matched with j under the new matching μ' . Condition (ii) states

that those unaffected by this change in partnership are matched to the same partner, under μ . Condition (iii) states that all other k , which would be the partners of i and j , if they were matched in μ , are single in μ' . This transition of μ to μ' is borrowed essentially from the process of "satisfying blocking pairs" in Roth and Vande Vate (1990), although we do not assume that (i, j) is a blocking pair so that they need not be better off in the matching μ' . Also, we allow for the possibility that $i = j$ in the definition, so that we cover the case in which only a single agent deviates.

For simplicity, we do not allow for simultaneous deviations by groups of three or more agents. Therefore, $\mu \rightarrow_S \mu'$ holds only if $S = \{i, j\}$ for some i, j . The results presented in this section occur when coalition deviations are allowed and defined as a suitable extension of the enforceability condition defined for pairs. See Mauleon et al. (2011) and Klaus et al. (2011) for details.

As in marriage games, the objective in roommate games is to form pairs. The difference, however, is that in roommate games, we do not have the set of players being partitioned into two disjoint sets, as in marriage games. In such sense, sometimes roommate games are called one-sided matching games, as opposed to the two sidedness in marriage games. When considering the mapping μ , for describing a matching, we do not have the restriction that $\mu(i)$ be a member of the opposite group if not matched to itself. However, this restriction is also lifted when considering blocking pairs, and this relative ease of forming a blocking pair contributes to the lack of general existence of a core matching in roommate games, as is shown in Gale and Shapley (1962).

The roommate game is given by the components (N, R) where N is the set of players, no longer partitioned into two disjoint sets, and each $i \in N$ has a preference relation R_i over N . As in the marriage game, a matching is a function $\mu: N \rightarrow N$ such that it is a bijection and satisfies the condition $\mu(i) = j$ if and only if $\mu(j) = i$. A pair $(i, j) \in N \times N$ is said to block μ if and only if $j P_i \mu(i)$ and $i P_j \mu(j)$ hold. We allow the case in which $i = j$ so that the previous definition also covers the case in which μ is blocked by a single person.

A matching μ is said to be a core matching if it is not blocked by any pair or single agent. A matching μ is said to be individually rational if it is not blocked by a single agent.

The indirect domination relation considered in both Mauleon et al. (2011) and Klaus et al. (2011) was defined in the usual way using the enforceability condition. Therefore, unlike the network formation model, we require in the indirect domination relation that when a pair deviates, both members must be better off in the final matching that is achieved in the sequence.

Mauleon et al. (2011) and Klaus et al. (2011) both proved the following lemma for the marriage game and roommate game, respectively. The lemma provides a relationship between indirect domination and blocking by pairs.

Lemma Let μ and μ' be individually rational matchings. A matching μ indirectly dominates μ' if and only if there does not exist a pair (i, j) that blocks μ such that $\mu'(i) = j$.

The previous lemma then leads to the connection between core matchings and farsighted stable sets of matching games and roommate games. We state the results of the two games in the same theorem below.

Theorem 23 The following hold for both the marriage game and the roommate game:

1. Let μ be a core matching, and then $\{\mu\}$ is a farsighted stable set.
2. If μ is a singleton farsighted stable set, then $\{\mu\}$ must be a core matching.

The first part establishes the existence of a farsighted stable set if a core matching exists – which is always satisfied in marriage games but not so in roommate games. The classic three-player example in which a core matching does not exist also is an example in which a farsighted stable set may not exist. The second statement says that every singleton farsighted stable set must arise from a core matching. This theorem by itself does not preclude the existence of farsighted stable sets. However, Mauleon et al.

(2011) proved that there are no other types of farsighted stable sets in marriage games. Klaus et al. (2011) showed through an example the possible existence of a farsighted stable set other than that described in the theorem, but they prove that there are no farsighted stable sets with exactly two matchings. In their example, a core matching does not exist, thus showing that a farsighted stable set can exist more often than core matchings.

The results for roommate markets are thus much less clearer, and it seems unlikely to give any general results on the existence of farsighted stable sets other than when a core matching exists. One direction that is taken in Mauleon et al. (2014) is to find out when the indirect domination relation and the domination relation coincide. If such a condition on preferences can be found, then a core matching would simultaneously be a matching that cannot be indirectly dominated.

House Barter Games

Next, we consider the house barter model of Shapley and Scarf (1974) which appeared in an earlier section. For this model, the enforceability relation is defined in a way that is similar to the coalition formation model. The following version, defined in Kawasaki (2010), was given in Klaus et al. (2010).

For an allocation x , draw a directed graph where the set of nodes is the set of players and there is a directed edge (i, j) if $x(i) = j$. Because x is a bijection on N , each i is in exactly one cycle of this graph which is called the trading cycle of i under allocation x . Denote by $C^{x,i}$ the unique trading cycle of allocation x that includes agent i . For any two allocations, x and y , the relation $x \rightarrow_S y$ holds if and only if the following three conditions are satisfied:

- $y(S) = S$
- $y(i) = x(i)$ if $i \notin \bigcup_{j \in S} C^{x,j}$
- $y(i) = i$ if $i \in \bigcup_{j \in S} C^{x,j}$ and $i \notin S$

The first condition states that the allocation y can be achieved by a reallocation of the endowments initially owned by members of the coalition S among themselves. The second condition states

that the allocation of the goods to those unaffected by the coalition S should be unchanged in the new allocation y . Finally, those who are affected are then assigned to their original endowment. This is not the only way to define the enforceability condition for these games. For another example, see Serrano and Volij (2008), although this is in a different framework.

We can now define indirect domination and direct domination for this framework. It should be noted that in the previous sections, we had treated direct domination as if it were equivalent to the myopic domination relations that have already been defined. This is not the case for the house barter model because of the details on the assignments to MS . For an example, see Kawasaki (2010).

An allocation x is said to be a competitive allocation if there exists a price system $(p_i)_{i \in N}$, where p_i denotes the price of the good initially owned by agent i , satisfying the following condition: $j \succ_i x_i \Rightarrow p_j > p_i$. Shapley and Scarf (1974) showed that the set of competitive allocations coincides with the allocations that can be obtained by the top trading cycle method.

Shapley and Scarf (1974) and Roth and Postlewaite (1977) showed that a competitive allocation exists, and every competitive allocation is in the core. Meanwhile, Wako (1984) showed that the strict core is a subset of the set of competitive allocations, and this inclusion can be strict. Thus, the set of competitive allocations can lie strictly between the strict core and the core.

Wako (1999) defined a domination relation, called antisymmetric weak domination (adom), such that the set of competitive allocations is the core defined by this domination relation. The formal definition is as follows.

An allocation x is said to antisymmetrically weakly dominate another allocation y if there exists a coalition S such that x weakly dominates y via S and $x(i) I_i y(i)$ holds only when $x(i) = y(i)$ and denote this relationship by $x \text{ adom } y$. The difference between adom and wdom is that in adom, the agents who find the two allocations indifferent do not receive a different good. By definition, adom is stronger than wdom but weaker than sdom. Therefore, the core defined

by adom should lie between the cores defined by sdom and wdom. Wako (1999) showed that the core defined by adom coincides with the set of competitive allocations. Furthermore, Toda (1997) showed that the core is the unique stable set defined by adom. Kawasaki (2010) showed that the analogue of these two results also holds for the respective sets defined by an indirect domination relation based on adom. The definition is as follows.

An allocation x is said to indirectly and antisymmetrically weakly dominate another allocation y and denote this by $x \text{ iadom } y$ if there exists a sequence of allocations $y = x^0, x^1, \dots, x^p = x$ and coalitions $S^1, S^2, \dots, S^p$ such that for each $j = 1, 2, \dots, p$, (i) $x^{j-1} \rightarrow_{S^j} x^j$, (ii) $x(i) R_i x^{j-1}(i)$ for all $i \in S^j$ with strict preference for some $i \in S^j$ and (iii) if for $i \in S^j$ $x^{j-1} I_i x$, then $x^{j-1}(i) = x^j(i)$.

The interpretation for the first two conditions is the same as in the indirect domination relation described in the earlier sections. The third condition incorporates the behavioral assumption in the adom relation into the iadom relation here. If an agent is indifferent between the good that is currently allocated and the one that will be allocated in the end, then that good that is allocated to that agent is unchanged in the reallocation described in (i). Kawasaki (2010) showed that the set of competitive allocations is the unique stable set defined by the iadom relation.

Because the iadom relation is quite complicated, Klaus et al. (2010) analyzed the stable sets using the usual indirect domination relation used in the literature built off of the enforceability condition in Kawasaki (2010). The definition of this indirect domination relation involves the existence of a sequence of allocations and coalitions satisfying (i) above, and conditions (ii) and (iii) are replaced by the condition $x(i) P_i x^{j-1}(i)$. Because this indirect domination is essentially the indirect domination in the Chwe's framework defined for this game, we will still use the term farsighted stable set to describe a stable set defined by this indirect domination relation.

Klaus et al. (2010) showed that Kawasaki's (2010) result can be obtained using indirect domination relation for markets that satisfy the following condition: for each $i \neq j$,

$$j P_i i \text{ or } i P_i j.$$

The condition states that for each agent, the good owned by other agents is either strictly better than his endowment or strictly worse. This condition is by no means restrictive and holds in many applications. Klaus et al. (2010) then proved the following result under this restriction on preferences and provide an example in which the result does not hold in a market that does not satisfy the above condition.

Theorem 24 The set of competitive allocations is the unique farsighted stable set (with respect to the usual indirect domination) for a house barter market satisfying the above condition.

Future Directions

We have reviewed applications of von Neumann-Morgenstern stable sets in several types of games pertaining to economic, political, and social systems. Stable sets give us insights into coalition formation among players in the systems in question. Farsighted stable sets, especially applied to some economic systems, show that players' farsighted behavior leads to Pareto-efficient outcomes even though their collaboration is not binding. The stable set analysis is also applicable to games with infinitely many players. Application to such models gives us new insights into behavior in large economic and social systems. For details, we refer the reader to Hart (1974), Einy et al. (1996), Einy and Shitovitz (1996), Greenberg et al. (1996), Shitovitz and Weber (1997), and Rosenmüller and Shitovitz (2000).

Another interesting model that has been studied but due to space constraints that have not been described in detail in this article includes pillage games of Jordan (2006). In pillage games, coalitions with more power can dictate the wealth of those with less power or *pillage* their wealth, and the concept of domination is then defined along those lines. It is a game form that cannot be expressed simply by a characteristic function. However, like some examples in characteristic

function form, there may exist a finite stable set even though the set of outcomes is an infinite set. Kerber and Rowat (2011) gives a bound as to the number of elements that can be included in such finite stable set by using graph theory.

Another strand of research that has been undertaken is one where farsighted concepts are incorporated into a more dynamic concept. (Connections between the classical stable set and some dynamic bargaining procedure in voting situations have been documented in papers by Anesi 2006, 2010; Diermeier and Fong 2012.) In the abstract model, each deviation in the sequence in the definition of indirect domination represented hypothetical moves each coalition can make, merely representing possible moves coalitions can make to better themselves in the long run. Another viewpoint is that each move in a sequence is an actual physical move resulting in actual payoffs in each step, but each coalition values the future sufficiently enough to focus on the end results. This approach is taken in Ray and Vohra (2015b), which is also reminiscent of models of coalition formation of Konishi and Ray (2003) and models of coalitional bargaining similar to that of Gomes (2005) and Gomes and Jehiel (2005). Dutta and Vohra (2017) incorporate some of the ideas in the dynamic framework back into the static framework to look more closely at the sequences in the definition of indirect domination. Their work presents a new way of refining explicitly the sequences in the indirect domination relation, an issue that have also been previously mentioned in papers such as Xue (1998) and Mauleon and Vannetelbosch (2004).

Finally, there have also been papers which have incorporated stable sets in environments with incomplete information. Graziano et al. (2015) define stable sets for economies with asymmetric information. There is also a study on the linkage between common knowledge of Bayesian rationality and achievement of stable sets in generalized abstract games. See Luo (2001, 2009) for details.

The realm of application of stable sets, both myopic and farsighted, can extend to games in which players inherently take both cooperative and noncooperative behaviors. Those studies

will in turn have impacts on developments of economics, politics, sociology, and many applied social sciences.

Acknowledgments This work was supported by the Japan Society for the Promotion of Science (JSPS) Grant Numbers JP16H03121 and JP17K13696.

Bibliography

- Anesi V (2006) Committee with farsighted voters: a new interpretation of stable sets. *Soc Choice Welf* 27:595–610
- Anesi V (2010) Noncooperative foundations of stable sets in voting games. *Games Econ Behav* 70:488–493
- Aumann R, Peleg B (1960) Von Neumann-Morgenstern solutions to cooperative games without side payments. *Bull Am Math Soc* 66:173–179
- Bala V, Goyal S (2000) A noncooperative model of network formation. *Econometrica* 68:1181–1229
- Bando K (2014) On the existence of a strictly strong Nash equilibrium under the student-optimal deferred acceptance algorithm. *Games Econ Behav* 87:269–287
- Banerjee S, Konishi H, Sonmez T (2001) Core in a simple coalition formation game. *Soc Choice Welf* 18:135–153
- Barberá S, Gerber A (2003) On coalition formation: durable coalition structures. *Math Soc Sci* 45:185–203
- Beal S, Durieu J, Solal P (2008) Farsighted coalitional stability in TU-games. *Math Soc Sci* 56:303–313
- Bernheim D, Peleg B, Whinston M (1987) Coalition-proof Nash equilibria: Concepts. *J Econ Theory* 42:1–12
- Bhattacharya A, Brosi V (2011) An existence result for farsighted stable sets of games in characteristic function form. *Int J Game Theory* 40:393–401
- Bogomolnaia A, Jackson MO (2002) The stability of hedonic coalition structures. *Games Econ Behav* 38:201–230
- Bott R (1953) Symmetric solutions to majority games. In: Kuhn HW, Tucker AW (eds) *Contribution to the theory of games, volume II*, Annals of mathematics studies, vol 28. Princeton University Press, Princeton, pp 319–323
- Chwe MS-Y (1994) Farsighted coalitional stability. *J Econ Theory* 63:299–325
- Diamantoudi E (2005) Stable cartels revisited. *Econ Theory* 26:907–921
- Diamantoudi E, Xue L (2003) Farsighted stability in hedonic games. *Soc Choice Welf* 21:39–61
- Diamantoudi E, Xue L (2007) Coalitions, agreements and efficiency. *J Econ Theory* 136:105–125
- Diermeier D, Fong P (2012) Characterization of the von Neumann-Morgenstern stable set in a non-cooperative model of dynamic policy-making with a persistent agenda setter. *Games Econ Behav* 76:349–353
- Dutta B, Vohra R (2017) Rational expectations and farsighted stability. *Theor Econ* 12:1191–1227
- Ehlers L (2007) Von Neumann-Morgenstern stable sets in matching problems. *J Econ Theory* 134:537–547
- Einy E, Shitovitz B (1996) Convex games and stable sets. *Games Econ Behav* 16:192–201
- Einy E, Holzman R, Monderer D, Shitovitz B (1996) Core and stable sets of large games arising in economics. *J Econ Theory* 68:200–211
- Funaki Y, Yamato T (2014) Stable coalition structures under restricted coalitional changes. *Int Game Theory Rev* 16. <https://doi.org/10.1142/S0219198914500066>
- Gale D, Shapley LS (1962) College admissions and the stability of marriage. *Am Math Mon* 69:9–15
- Gomes A (2005) Multilateral contracting with externalities. *Econometrica* 73:1329–1350
- Gomes A, Jehiel P (2005) Dynamic process of social and economic interactions: on the persistence of inefficiencies. *J Pol Econ* 113:626–667
- Graziano MG, Meo C, Yannelis NC (2015) Stable sets for asymmetric information economies. *Int J Econ Theory* 11:137–154
- Greenberg J (1990) *The theory of social situations: an alternative game theoretic approach*. Cambridge University Press, Cambridge
- Greenberg J, Monderer D, Shitovitz B (1996) Multistage situations. *Econometrica* 64:1415–1437
- Greenberg J, Luo X, Oladi R, Shitovitz B (2002) (Sophisticated) stable sets in exchange economies. *Games Econ Behav* 39:54–70
- Griesmer JH (1959) Extreme games with three values. In: Tucker AW, Luce RD (eds) *Contribution to the theory of games, vol IV*. Annals of mathematics studies, vol 40. Princeton University Press, Princeton, pp 189–212
- Gusfield D, Irving R (1989) *The stable marriage problem: structure and algorithms*. MIT Press, Boston
- Harsanyi J (1974) An equilibrium-point interpretation of stable sets and a proposed alternative definition. *Manag Sci* 20:1472–1495
- Hart S (1973) Symmetric solutions of some production economies. *Int J Game Theory* 2:53–62
- Hart S (1974) Formation of cartels in large markets. *J Econ Theory* 7:453–466
- Heijmans J (1991) Discriminatory von Neumann-Morgenstern solutions. *Games Econ Behav* 3:438–452
- Herings PJJ, Mauleon A, Vannetelbosch V (2009) Farsightedly stable networks. *Games Econ Behav* 67:526–541
- Herings PJJ, Mauleon A, Vannetelbosch V (2010) Coalition formation among farsighted agents. *Games* 1:286–298
- Herings PJJ, Mauleon A, Vannetelbosch V (2017) Stable sets in matching problems with coalitional sovereignty and path dominance. *J Math Econ* 71:14–19
- Herings PJJ, Mauleon A, Vannetelbosch V (2018) Stability of networks under horizon- K farsightedness. *Econ Theory*. <https://doi.org/10.1007/s00199-018-1119-7>
- Hirai T (2017) Single payoff farsighted stable sets in strategic games with punishment strategies. *Int J Game Theory*. <https://doi.org/10.1007/s00182-017-0597-3>
- Jackson MO, van den Nouweland A (2005) Strongly stable networks. *Games Econ Behav* 51:420–444

- Jackson MO, Wolinsky A (1996) A strategic model of social and economic networks. *J Econ Theory* 71:44–74
- Jordan JS (2006) Pillage and property. *J Econ Theory* 131:26–44
- Kamijo Y, Muto S (2010) Farsighted coalitional stability of a price leadership cartel. *Jpn Econ Rev* 61:455–465
- Kaneko M (1987) The conventionally stable sets in non-cooperative games with limited observations I: definition and introductory argument. *Math Soc Sci* 13:93–128
- Kawasaki R (2010) Farsighted stability of the competitive allocations in an exchange economy with indivisible goods. *Math Soc Sci* 59:46–52
- Kawasaki R (2015) Maximin, minimax, and von Neumann-Morgenstern farsighted stable sets. *Math Soc Sci* 74:8–12
- Kawasaki R, Muto S (2009) Farsighted stability in provision of perfectly lumpy public goods. *Math Soc Sci* 58:98–109
- Kawasaki R, Sato T, Muto S (2015) Farsightedly stable tariffs. *Math Soc Sci* 76:118–124
- Kerber M, Rowat C (2011) A Ramsey bound on stable sets in Jordan pillage games. *Int J Game Theory* 40:461–466
- Kirchsteiger G, Mantovani M, Mauleon A, Vannetelbosch V (2016) Limited farsightedness in network formation. *J Econ Behav Organ* 128:97–120
- Konishi H, Ray D (2003) Coalition formation as a dynamic process. *J Econ Theory* 110:1–41
- Klaus B, Klijn F, Walzl M (2010) Farsighted house allocation. *J Math Econ* 46:817–824
- Klaus B, Klijn F, Walzl M (2011) Farsighted stability for roommate markets. *J Pub Econ Theory* 13:921–933
- Lucas WF (1968) A game with no solution. *Bull Am Math Soc* 74:237–239
- Lucas WF (1990) Developments in stable set theory. In: Ichiiishi T et al (eds) *Game theory and applications*. Academic, New York, pp 300–316
- Lucas WF, Rabie M (1982) Games with no solutions and empty cores. *Math Oper Res* 7:491–500
- Lucas WF, Michaelis K, Muto S, Rabie M (1982) A new family of finite solutions. *Int J Game Theory* 11:117–127
- Lucas WF (1992) Von Neumann-Morgenstern stable sets. In: Aumann RJ, Hart S (eds) *Handbook of game theory with economic applications*, vol 1. North-Holland, Amsterdam, pp 543–590
- Luo X (2001) General systems and ϕ -stable sets: a formal analysis of socioeconomic environments. *J Math Econ* 36:95–109
- Luo X (2009) On the foundation of stability. *Econ Theory* 40:185–201
- Manlove D (2013) *Algorithmics of matching under preferences*. World Scientific, Singapore
- Mariotti M (1997) A model of agreements in strategic form games. *J Econ Theory* 74:196–217
- Masuda T (2002) Farsighted stability in average return games. *Math Soc Sci* 44:169–181
- Mauleon A, Vannetelbosch V (2004) Farsighted and cautiousness in coalition formation games with positive spillovers. *Theory Dec* 56:291–324
- Mauleon A, Vannetelbosch V, Vergote W (2011) Von Neumann-Morgenstern farsightedly stable sets in two-sided matching. *Theor Econ* 6:499–521
- Mauleon A, Molis E, Vannetelbosch V, Vergote W (2014) Dominance invariant one-to-one matching problems. *Int J Game Theory* 43:925–943
- Moulin H (1995) *Cooperative microeconomics: a game-theoretic introduction*. Princeton University Press, Princeton
- Muto S (1979) Symmetric solutions for symmetric constant-sum extreme games with four values. *Int J Game Theory* 8:115–123
- Muto S (1982a) On Hart production games. *Math Oper Res* 7:319–333
- Muto S (1982b) Symmetric solutions for (n, k) games. *Int J Game Theory* 11:195–201
- Muto S, Okada D (1996) Von Neumann-Morgenstern stable sets in a price-setting duopoly. *Econ Econ* 81:1–14
- Muto S, Okada D (1998) Von Neumann-Morgenstern stable sets in Cournot competition. *Econ Econ* 85:37–57
- Nakanishi N (1999) Reexamination of the international export quota game through the theory of social situations. *Games Econ Behav* 27:132–152
- Nakanishi N (2001) On the existence and efficiency of the von Neumann-Morgenstern stable set in an n -player prisoner's dilemma. *Int J Game Theory* 30:291–307
- Nakanishi N (2009) Noncooperative farsighted stable set in an n -player prisoners' dilemma. *Int J Game Theory* 38:249–261
- Nakayama M (1998) Self-binding coalitions. *Keio Econ Stud* 35:1–8
- Núñez M, Rafels C (2013) Von Neumann-Morgenstern solutions in the assignment market. *J Econ Theory* 148:1282–1291
- Oladi R (2005) Stable tariffs and retaliation. *Rev Int Econ* 13:205–215
- Owen G (1965) A class of discriminatory solutions to simple N -person games. *Duke Math J* 32:545–553
- Owen G (1968) n -Person games with only 1, $n-1$, and n -person coalitions. *Proc Am Math Soc* 19:1258–1261
- Owen G (1995) *Game theory*, 3rd edn. Academic, New York
- Page FH, Wooders M (2009) Strategic basins of attraction, the path dominance core, and network formation games. *Games Econ Behav* 66:462–487
- Page FH, Wooders MH, Kamat S (2005) Networks and farsighted stability. *J Econ Theory* 120:257–269
- Peleg B (1986) A proof that the core of an ordinal convex game is a von Neumann-Morgenstern solution. *Math Soc Sci* 11:83–87
- Quint T, Wako J (2004) On houseswapping, the strict core, segmentation, and linear programming. *Math Oper Res* 29:861–877
- Ray D, Vohra R (1997) Equilibrium binding agreements. *J Econ Theory* 73:30–78
- Ray D, Vohra R (2015a) The farsighted stable set. *Econometrica* 83:977–1011
- Ray D, Vohra R (2015b) Coalition formation. In: Young HP, Zamir S (eds) *Handbook of game theory*, vol 4. Elsevier/North Holland, pp 239–326
- Rosenmüller J (1977) *Extreme games and their solutions*, Lecture notes in economics and mathematical systems, vol 145. Springer, Berlin

- Rosenmüller J, Shitovitz B (2000) A characterization of vNM-stable sets for linear production games. *Int J Game Theory* 29:39–61
- Roth AE, Postlewaite A (1977) Weak versus strong domination in a market with indivisible goods. *J Math Econ* 4:131–137
- Roth AE, Sotomayor MO (1990) Two-sided matching: a study in game-theoretic modeling and analysis. Cambridge University Press, Cambridge
- Roth AE, Vande Vate JH (1990) Random paths to stability in two-sided matching. *Econometrica* 58:1475–1480
- Serrano R, Volij O (2008) Mistakes in cooperation: the stochastic stability of Edgeworth's recontracting. *Econ J* 118:1719–1741
- Shapley LS (1953) Quota solutions of n -person games. In: Kuhn HW, Tucker TW (eds) *Contribution to the theory of games*, vol II. *Annals of mathematics studies*, vol 28. Princeton University Press, Princeton, pp 343–359
- Shapley LS (1959) The solutions of a symmetric market game. In: Tucker AW, Luce RD (eds) *Contribution to the theory of games*, vol IV. *Annals of mathematics studies*, vol 40. Princeton University Press, Princeton, pp 145–162
- Shapley LS (1962) Simple games: an outline of the descriptive theory. *Behav Sci* 7:59–66
- Shapley LS (1964) Solutions of compound simple games. In: Tucker AW et al (eds) *Advances in game theory*. *Annals of mathematics studies*, vol 52. Princeton University Press, Princeton, pp 267–305
- Shapley LS (1971) Cores of convex games. *Int J Game Theory* 1:11–26
- Shapley LS, Scarf H (1974) On cores and indivisibilities. *J Math Econ* 1:23–37
- Shapley LS, Shubik M (1972) The assignment game I: the core. *Int J Game Theory* 1:111–130
- Shino J, Kawasaki R (2012) Farsighted stable sets in Hotelling's location games. *Math Soc Sci* 63:23–30
- Shitovitz B, Weber S (1997) The graph of Lindahl correspondence as the unique von Neumann-Morgenstern abstract stable set. *J Math Econ* 27:375–387
- Shubik M (1982) *Game theory in the social sciences: concepts and solutions*. MIT Press, Boston
- Shubik M (1985) *A game-theoretic approach to political economy*. MIT Press, Boston
- Simonard M (1966) *Linear programming*. Prentice-Hall, Englewood Cliffs
- Solymosi T, Raghavan TES (2001) Assignment games with stable core. *Int J Game Theory* 30:177–185
- Sung SC, Dimitrov D (2007) On myopic stability concepts for hedonic games. *Theory Dec* 62:31–45
- Suzuki A, Muto S (2000) Farsighted stability in prisoner's dilemma. *J Oper Res Soc Jpn* 43:249–265
- Suzuki A, Muto S (2005) Farsighted stability in n -person prisoner's dilemma. *Int J Game Theory* 33:431–445
- Suzuki A, Muto S (2006) Farsighted behavior leads to efficiency in duopoly markets. In: Haurie A et al (eds) *Advances in dynamic games*. Birkhauser, Boston, pp 379–395
- Toda M (1997) Implementation and characterizations of the competitive solution with indivisibility. *Mimeo*
- von Neumann J, Morgenstern O (1953) *Theory of games and economic behavior*, 3rd edn. Princeton University Press, Princeton
- Wako J (1984) A note on the strong core of a market with indivisible goods. *J Math Econ* 13:189–194
- Wako J (1991) Some properties of weak domination in an exchange market with indivisible goods. *Jpn Econ Rev* 42:303–314
- Wako J (1999) Coalitional-proofness of the competitive allocations in an indivisible goods market. *Fields Inst Commun* 23:277–283
- Wako J (2010) A polynomial-time algorithm to find von Neumann-Morgenstern stable matchings in marriage games. *Algorithmica* 58:188–220
- Xue L (1997) Nonemptiness of the largest consistent set. *J Econ Theory* 73:453–459
- Xue L (1998) Coalitional stability under perfect foresight. *Econ Theory* 11:603–627
- Zhang J, Xue L, Zu L (2013) Farsighted free trade networks. *Int J Game Theory* 42:375–398

Cooperative Games

Roberto Serrano
Department of Economics, Brown University,
Providence, RI, USA

Article Outline

Glossary
Definition of the Subject
Introduction
Cooperative Games
The Core
The Shapley Value
Future Directions
Bibliography

Core

feasible payoffs for each coalition, what payoff will be awarded to each player? One can take a positive or normative approach to answering this question, and different solution concepts in the theory lean towards one or the other.

It is a solution concept that assigns to each cooperative game the set of payoffs that no coalition can improve upon or block. In a context in which there is unfettered coalitional interaction, the core arises as a good positive answer to the question posed in cooperative game theory. In other words, if a payoff does not belong to the core, one should not expect to see it as the prediction of the theory if there is full cooperation. It is a solution that prescribes a single payoff for each player, which is the average of all marginal contributions of that player to each coalition he or she is a member of. It is usually viewed as a good normative answer to the question posed in cooperative game theory. That is, those who contribute more to the groups that include them should be paid more.

Glossary

Characteristic or coalitional function The most usual way to represent a cooperative game.

Cooperative game Strategic situation involving coalitions, whose formation assumes the existence of binding agreements among players.

Core Solution concept that assigns the set of payoffs that cannot be improved upon by any coalition.

Game theory Discipline that studies strategic situations.

Shapley value Solution concept that assigns the average of marginal contributions to coalitions.

Solution concept Mapping that assigns predictions to each game.

Definition of the Subject

Cooperative It is one of the two counterparts of game theory. It studies the interactions among coalitions of game theory players. Its main question is this: Given the sets of

Although there were some earlier contributions, the official date of birth of game theory is usually taken to be 1944, year of publication of the first edition of the *Theory of Games and Economic Behavior*, by John von Neumann and Oskar Morgenstern (1944). The core was first proposed by Francis Ysidro Edgeworth in 1881 (Edgeworth 1881), and later reinvented and defined in game theoretic terms in Gillies (1959). The Shapley value was proposed by Lloyd Shapley in his 1953

Ph.D. dissertation (Shapley 1953). Both the core and the Shapley value have been applied widely, to shed light on problems in different disciplines, including economics and political science.

Introduction

Game theory is the study of games, also called strategic situations. These are decision problems with multiple decision makers, whose decisions impact one another. It is divided into two branches: non-cooperative game theory and cooperative game theory. The actors in non-cooperative game theory are individual players, who may reach agreements only if they are self-enforcing. The non-cooperative approach provides a rich language and develops useful tools to analyze games. One clear advantage of the approach is that it is able to model how specific details of the interaction among individual players may impact the final outcome. One limitation, however, is that its predictions may be highly sensitive to those details. For this reason it is worth also analyzing more abstract approaches that attempt to obtain conclusions that are independent of such details. The cooperative approach is one such attempt, and it is the subject of this article.

The actors in cooperative game theory are coalitions, that is, groups of players. For the most part, two facts, that coalitions can form and that each coalition has a feasible set of payoffs available to its members, are taken as given the coalitions and their sets of feasible payoffs as primitives, the question tackled is the identification of final payoffs awarded to each player. That is, given a collection of feasible sets of payoffs, one for each coalition, can one predict or recommend a payoff (or set of payoffs) to be awarded to each player? Such predictions or recommendations are embodied in different solution concepts.

Indeed, one can take several approaches to answering the question just posed. From a positive or descriptive point of view, one may want to get a prediction of the likely outcome of the interaction among the players, and hence, the resulting payoff be understood as the natural consequence of the forces at work in the system. Alternatively,

one can take a normative or prescriptive approach, set up a number of normative goals, typically embodied in axioms, and try to derive their logical implications. Although authors sometimes disagree on the classification of the different solution concepts according to these two criteria – as we shall see, the understanding of each solution concept is enhanced if one can view it from very distinct approaches – in this article we shall exemplify the positive approach with the core and the normative approach with the Shapley value. While this may oversimplify the issues, it should be helpful to a reader new to the subject.

The rest of the article is organized as follows. Section “[Cooperative Games](#)” introduces the basic model of a cooperative game, and discusses its assumptions as well as the notion of solution concepts. Section “[The Core](#)” is devoted to the core, and section “[The Shapley Value](#)” to the Shapley value. In each case, some of the main results for each of the two are described, and examples are provided. Section “[Future Directions](#)” discusses some directions for future research.

Cooperative Games

Representations of Games. The Characteristic Function

Let us begin by presenting the different ways to describe a game. The first two are the usual ways employed in non-cooperative game theory.

The most informative way to describe a game is called its *extensive form*. It consists of a game tree, specifying the timing of moves for each player and the information available to each of them at the time of making a move. At the end of each path of moves, a final outcome is reached and a payoff vector is specified. For each player, one can define a *strategy*, i.e., a complete contingent plan of action to play the game. That is, a strategy is a function that specifies a feasible move each time a player is called upon to make a move in the game.

One can abstract from details of the interaction (such as timing of moves and information available at each move), and focus on the concept of strategies. That is, one can list down the set of strategies available to each player, and arrive at the *strategic*

or *normal form* of the game. For two players, for example, the normal form is represented in a bimatrix table. One player controls the rows, and the other the columns. Each cell of the bimatrix is occupied with an ordered pair, specifying the payoff to each player if each of them chooses the strategy corresponding to that cell.

One can further abstract from the notion of strategies, which will lead to the *characteristic function form* of representing a game. From the strategic form, one makes assumptions about the strategies used by the complement of a coalition of players to determine the feasible payoffs for the coalition (see, for example, the derivations in Aumann and Peleg (1960), von Neumann and Morgenstern (1944)). This is the representation most often used in cooperative game theory.

Thus, here are the primitives of the basic model in cooperative game theory. Let $N = \{1, \dots, n\}$ be a finite set of players. Each non-empty subset of N is called a *coalition*. The set N is referred to as the *grand coalition*. For each coalition S , we shall specify a set $V(S) \subset \mathbb{R}^{|S|}$ containing $|S|$ -dimensional payoff vectors that are feasible for coalition S . This is called the *characteristic function*, and the pair (N, V) is called a *cooperative game*. Note how a reduced form approach is taken because one does not explain what strategic choices are behind each of the payoff vectors in $V(S)$. In addition, in this formulation, it is implicitly assumed that the actions taken by the complement coalition (those players in $N \setminus S$) cannot prevent S from achieving each of the payoff vectors in $V(S)$. There are more general models in which these sorts of externalities across coalitions are considered, but we shall ignore them in this article.

Assumptions on the Characteristic Function

Some of the most common technical assumptions made on the characteristic function are the following:

1. For each $S \subseteq N$, $V(S)$ is closed. Denote by $\partial V(S)$ the boundary of $V(S)$. Hence, $\partial V(S) \subseteq V(S)$.
2. For each $S \subseteq N$, $V(S)$ is comprehensive, i.e., for each $x \in V(S)$, $\{x\} - \mathbb{R}_+^{|S|} \subseteq V(S)$.

3. For each $x \in \mathbb{R}^{|S|}$,

$$\partial V(S) \cap \left(\{x\} + \mathbb{R}_+^{|S|} \right)$$

is bounded.

4. For each $S \subseteq N$, there exists a continuously differentiable representation of $V(S)$, i.e., a continuously differentiable function $g_S: \mathbb{R}^{|S|} \rightarrow \mathbb{R}$ such that

$$V(S) = \{x \in \mathbb{R}^{|S|} : g_S(x) \leq 0\}.$$

5. For each $S \subseteq N$, $V(S)$ is non-leveled, i.e., for every $x \in \partial V(S)$ the gradient of g_S at x is positive in all its coordinates.

With the assumptions made, $\partial V(S)$ is its Pareto frontier, i.e., the set of vectors $x_S \in V(S)$ such that there does not exist $y_S \in V(S)$ satisfying that $y_i \geq x_i$ for all $i \in S$ with at least one strict inequality. Other assumptions usually made relate the possibilities available to different coalitions. Among them, a very important one is *balancedness*, which we shall define next:

A collection $\mathcal{T}$ of coalitions is balanced if there exists a set of weights $w(S) \in [0, 1]$ for each $S \in \mathcal{T}$ such that for every $i \in N$, $\sum_{S \in \mathcal{T}, i \in S} w(S) = 1$. One can think of these weights as the fraction of time that each player devotes to each coalition he is a member of, with a given coalition representing the same fraction of time for each player. The game (N, V) is balanced if $x_N \in V(N)$ whenever $(x_S) \in V(S)$ for every S in a balanced collection $\mathcal{T}$. That is, the grand coalition can always implement any “time-sharing arrangement” that the different sub-coalitions may come up with.

The characteristic function defined so far is often referred to as a *non-transferable utility (NTU) game*. A particular case is the *transferable utility (TU) game* case, in which for each coalition $S \subseteq N$, there exists a real number $v(S)$ such that

$$V(S) = \left\{ x \in \mathbb{R}^{|S|} : \sum_{i \in S} x_i \leq v(S) \right\}.$$

Abusing notation slightly, we shall denote a TU game by (N, v) . In the TU case there is an underlying

numeraire – money – that can transfer utility or payoff at a one-to-one rate from one player to any other. Technically, the theory of NTU games is far more complex: it uses convex analysis and fixed point theorems, whereas the TU theory is based on linear inequalities and combinatorics.

Solution Concepts

Given a characteristic function, i.e., a collection of sets $V(S)$, one for each S , the theory formulates its predictions on the basis of different *solution concepts*. We shall concentrate on the case in which the grand coalition forms, that is, cooperation is totally successful. Of course, solution concepts can be adapted to take care of the case in which this does not happen.

A *solution* is a mapping that assigns a set of payoff vectors in $V(N)$ to each characteristic function game (N, V) . Thus, a solution in general prescribes a set, which can be empty, or a singleton (when it assigns a unique payoff vector as a fraction of the fundamentals of the problem). The leading set-valued cooperative solution concept is the core, while one of the most used single-valued ones is the Shapley value for TU games.

There are several criteria to evaluate the reasonableness or appeal of a cooperative solution. As outlined above, in a normative approach, one can propose axioms, abstract principles that one would like the solution to satisfy, and the next step is to pursue their logical consequences. Historically, this was the first argument to justify the Shapley value. Alternatively, one could start by defending a solution on the basis of its definition alone. In the case of the core, this will be especially natural: in a context in which players can freely get together in groups, the prediction should be payoff vectors that cannot be improved upon by any coalition. One can further enhance one's positive understanding of the solution concept by proposing games in extensive form or in normal form played non-cooperatively by players whose self-enforcing agreements lead to a given solution. This is simply to provide non-cooperative foundations or non-cooperative implementation to the cooperative solution in question, and it is an important research agenda initiated by John Nash in (Nash 1953), referred to as the Nash program (see Serrano

(2005) for a recent survey). Today, there are interesting results of these different kinds for many solution concepts, which include axiomatic characterizations and non-cooperative foundations. Thus, one can evaluate the appeal of the axioms and the non-cooperative procedures behind each solution to defend a more normative or positive interpretation in each case.

The Core

The idea of agreements that are immune to coalitional deviations was first introduced to economic theory by Edgeworth in (Edgeworth 1881), which defined the set of coalitionally stable allocations of an economy under the name “final settlements.” Edgeworth envisioned this concept as an alternative to *competitive equilibrium* (Walras 1874), of central importance in economic theory, and was also the first to investigate the connections between the two concepts. Edgeworth's notion, which today we refer to as the *core*, was rediscovered and introduced to game theory in Gillies (1959). The origins of the core were not axiomatic. Rather, its simple and appealing definition appropriately describes stable outcomes in a context of unfettered coalitional interaction.

The core of the game (N, V) is the set of payoff vectors

$$C(N, V) = \{x \in V(N) : \nexists S \subseteq N, x_S \in V(S) \setminus \partial V(S)\}.$$

In words, it is the set of feasible payoff vectors for the grand coalition that no coalition can upset. If such a coalition S exists, we shall say that S can improve upon or block x , and x is deemed unstable. That is, in a context where any coalition can get together, when S has a blocking move, coalition S will form and abandon the grand coalition and its payoffs x_S in order to get to a better payoff for each of the members of the coalition, a plan that is feasible for them.

Non-Emptiness

The core can prescribe the empty set in some games. A game with an empty core is to be understood as a situation of strong instability, as any

payoffs proposed to the grand coalition are vulnerable to coalitional blocking.

Example Consider the following simple majority 3-player TU game, in which the votes of at least two players makes the coalition winning. That is, we represent the situation by the following characteristic function: $v(S) = 1$ for any S containing at least two members, $v(\{i\}) = 0$ for all $i \in N$. Clearly, $C(N, v) = \emptyset$. Any feasible payoff agreement proposed to the grand coalition will be blocked by at least one coalition.

An important sufficient condition for the non-emptiness of the core of NTU games is balancedness, as shown in Scarf (1967):

Theorem 1 (Scarf (1967)) Let the game (N, V) be balanced. Then $C(N, V) \neq \emptyset$.

For the TU case, balancedness is not only sufficient, but it becomes also necessary for the non-emptiness of the core:

Theorem 2 (Bondareva (1963); Shapley (1967)) Let (N, v) be a TU game. Then, (N, v) is balanced if and only of $C(N, V) \neq \emptyset$.

The Connections with Competitive Equilibrium

In economics, the institution of markets and the notion of prices are essential to the understanding of the allocation of goods and the distribution of wealth among individuals. For simplicity in the presentation, we shall concentrate on exchange economies, and disregard production aspects. That is, we shall assume that the goods in question have already been produced in some fixed amounts, and now they are to be allocated to individuals to satisfy their consumption needs.

An *exchange economy* is a system in which each agent i in the set N has a consumption set $Z_i \subseteq \mathbb{R}_+^l$ of commodity bundles, as well as a preference relation over Z_i and an initial endowment $\omega_i \in Z_i$ of the commodities. A feasible allocation of goods in the economy is a list of bundles $(Z_i)_{i \in N}$ such that $Z_i \in Z_i$ and $\sum_{i \in N} z_i \leq \sum_{i \in N} \omega_i$. An allocation is *competitive* if it is supported by a *competitive equilibrium*. A competitive equilibrium is a price-allocation pair $(p, (z_i)_{i \in N})$, where $p \in \mathbb{R}^l \setminus \{0\}$ is such that

- for every $i \in N$, z_i is top-ranked for agent i among all bundles z satisfying that $pz \leq p\omega_i$,
- and $\sum_{i \in N} z_i = \sum_{i \in N} \omega_i$

In words, this is what the concept expresses. First, at the equilibrium prices, each agent demands z_i , i.e., wishes to purchase this bundle among the set of affordable bundles, the budget set. And second, these demands are such that all markets clear, i.e., total demand equals total supply.

Note how the notion of a competitive equilibrium relies on the principle of private ownership (each individual owns his or her endowment, which allows him or her to access markets and purchase things). Moreover, each agent is a price-taker in all markets. That is, no single individual can affect the market prices with his or her actions; prices are fixed parameters in each individual's consumption decision. The usual justification for the price-taking assumption is that each individual is "very small" with respect to the size of the economy, and hence, has no market power.

One difficulty with the competitive equilibrium concept is that it does not explain where prices come from. There is no single agent in the model responsible for coming up with them. Walras in (Walras 1874) told the story of an auctioneer calling out prices until demand and supply coincide, but in many real-world markets there is no auctioneer. More generally, economists attribute the equilibrium prices to the workings of the forces of demand and supply, but this appears to be simply repeating the definition. So, is there a different way one can explain competitive equilibrium prices?

As it turns out, there is a very robust result that answers this question. We refer to it as the *equivalence principle* (see, e.g., Aumann 1987), by which, under certain regularity conditions, the predictions provided by different game-theoretic solution concepts, when applied to an economy with a large enough set of agents, tend to converge to the set of competitive equilibrium allocations. One of the first results in this tradition was provided by Edgeworth in 1881 for the core. Note how the core of the economy can be defined in the space of allocations, using the same definition as

above. Namely, a feasible allocation is in the core if it cannot be blocked by any coalition of agents when making use of the coalition's endowments.

Edgeworth's result was generalized later by Debreu and Scarf in (Debreu and Scarf 1963) for the case in which an exchange economy is replicated an arbitrary number of times (Anderson studies in (Anderson 1978) the more general case of arbitrary sequences of economies, not necessarily replicas). An informal statement of the Debreu-Scarf theorem follows:

Theorem 3 (Debreu and Scarf (1963)) Consider an exchange economy. Then,

1. The set of competitive equilibrium allocations is contained in the core.
2. For each non-competitive core allocation of the original economy, there exists a sufficiently large replica of the economy for which the replica of the allocation is blocked.

The first part states a very appealing property of competitive allocations, i.e., their coalitional stability. The second part, known as the core convergence theorem, states that the core “shrinks” to the set of competitive allocations as the economy grows large.

In Aumann (1964), Aumann models the economy as an atomless measure space, and demonstrates the following core equivalence theorem:

Theorem 4 (Aumann (1964)) Let the economy consists of an atomless continuum of agents. Then, the core coincides with the set of competitive allocations.

For readers who wish to pursue the topic further, (Anderson 2008) provides a recent survey.

Axiomatic Characterizations

The axiomatic foundations of the core were provided much later than the concept was proposed. These characterizations are all inspired by Peleg's work. They include (Peleg 1985, 1986), and (Serrano and Volij 1998) – the latter paper also provides an axiomatization of competitive allocations in which core convergence insights are exploited.

In all these characterizations, the key axiom is that of *consistency*, also referred to as the reduced game property. Consistency means that the outcomes prescribed by a solution should be “invariant” to the number of players in the game. More formally, let (N, V) be a game, and let σ be a solution. Let $x \in \sigma(N, V)$. Then, the solution is consistent if for every $S \subseteq N$, $x_S \in \sigma(S, V_{x_S})$ where (S, V_{x_S}) is the reduced game for S given payoffs x , defined as follows. The feasible set for S in this reduced game is the projection of $V(N)$ at $x_{N \setminus S}$ i.e., remains after paying those outside of S :

$$V_{x_S}(S) = \{y_S : (y_S, x_{N \setminus S}) \in V(N)\}.$$

However, the feasible set of $T \subset S, T \neq S$, allows T to make deals with any coalition outside of S , provided that those services are paid at the rate prescribed by $x_{N \setminus S}$:

$$V_{x_S}(T) = \{y_T \in \cup_{Q \subseteq N \setminus S} (y_T, x_Q) \in V(T \cup Q)\}.$$

It can be shown that the core satisfies consistency with respect to this reduced game. Moreover, consistency is the central axiom in the characterization of the core, which, depending on the version one looks at, uses a host of other axioms; see Peleg (1985, 1986), Serrano and Volij (1998).

Non-cooperative Implementation

To obtain a non-cooperative implementation of the core, the procedure must embody some feature of anonymity, since the core is usually a large set and it contains payoffs where different players are treated very differently. For instance, if the procedure always had a fixed set of moves, typically the prediction would favor the first mover, making it impossible to obtain an implementation of the entire set of payoffs.

The model in Perry and Reny (1994) builds in this anonymity by assuming that negotiations take place in continuous time, so that anyone can speak at the beginning of the game, and at any point in time, instead of having a fixed order. The player that gets to speak first makes a proposal consisting of naming a coalition that contains him and a

feasible payoff for that coalition. Next, the players in that coalition get to respond. If they all accept the proposal, the coalition leaves and the game continues among the other players. Otherwise, a new proposal may come from any player in N . It is shown that, if the TU game has a non-empty core (as well as any of its subgames), a class of stationary self-enforcing predictions of this procedure coincide with the core. If a core payoff is proposed to the grand coalition, there are no incentives for individual players to reject it. Conversely, a non-core payoff cannot be sustained because any player in a blocking coalition has an incentive to make a proposal to that coalition, who will accept it (knowing that the alternative, given stationarity, would be to go back to the non-core status quo). Moldovanu and Winter (1995) offers a discrete-time version of the mechanism: in this work, the anonymity required is imposed on the solution concept, by looking at the order-independent equilibria of the procedure.

The model in Serrano (1995) sets up a market to implement the core. The anonymity of the procedure stems from the random choice of broker. The broker announces a vector $(x_1, \dots, x_n)$, where the components add up to $v(N)$. One can interpret x_i as the price for the productive asset held by player i . Following an arbitrary order, the remaining players either accept or reject these prices. If player i accepts, he sells his asset to the broker for the price x_i and leaves the game. Those who reject get to buy from the broker, at the called out prices, the portfolio of assets of their choice if the broker still has them. If a player rejects, but does not get to buy the portfolio of assets he would like because someone else took them before, he can always leave the market with his own asset. The broker's payoff is the worth of the final portfolio of assets that he holds, plus the net monetary transfers that he has received. It is shown in Serrano (1995) that the prices announced by the broker will always be his top-ranked vectors in the core. If the TU game is such that gains from cooperation increase with the size of coalitions, a beautiful theorem of Shapley in (Shapley 1971) is used to prove that the set of all equilibrium payoffs of this procedure will coincide with the core. Core payoffs are here understood as those price

vectors where all arbitrage opportunities in the market have been wiped out. Also, procedures in Serrano and Vohra (1997) implement the core, but do not rely on the TU assumption, and they use a procedure in which the order of moves can be endogenously changed by players. Finally, yet another way to build anonymity in the procedure is by allowing the proposal to be made by brokers outside of the set N , as done in Pérez-Castrillo (1994).

An Application

Consider majority games within a parliament. Suppose there are 100 seats, and decisions are made by simple majority so that 51 votes are required to pass a piece of legislation.

In the first specification, suppose there is a very large party – player 1 –, who has 90 seats. There are five small parties, with 2 seats each. Given the simple majority rules, this problem can be represented by the following TU characteristic function: $v(S) = 1$ if S contains player 1, and $v(S) = 0$ otherwise. The interpretation is that each winning coalition can get the entire surplus – pass the desired proposal. Here, a coalition is winning if and only if player 1 is in it. For this problem, the core is a singleton: the entire unit of surplus is allocated to player 1, who has all the power. Any split of the unit surplus of the grand coalition ($v(N) = 1$) that gives some positive fraction of surplus to any of the small parties can be blocked by the coalition of player 1 alone.

Consider now a second problem, in which player 1, who continues to be the large party, has 35 seats, and each of the other five parties has 13 seats. Now, the characteristic function is as follows: $v(S) = 1$ if and only if S either contains player 1 and two small parties, or it contains four of the small parties; $v(S) = 0$ otherwise. It is easy to see that now the core is empty: any split of the unit surplus will be blocked by at least one coalition. For example, the entire unit going to player 1 is blocked by the coalition of all five small parties, which can award 0.2 to each of them. But this arrangement, in which each small party gets 0.2 and player 1 nothing, is blocked as well, because player 1 can bribe two of the small parties (say, players 2 and 3) and promise them $1/3$ each,

keeping the other third for itself, and so on. The emptiness of the core is a way to describe the fragility of any agreement, due to the inherent instability of this coalition formation game.

The Shapley Value

Now consider a transferable utility or TU game in characteristic function form. The number $v(S)$ is referred to as the worth of S , and it expresses S 's initial position (e.g., the maximum total amount of surplus in numeraire -money, or power – that S initially has at its disposal).

Axiomatics

Shapley in (Shapley 1953) is interested in solving in a fair and unique way the problem of distribution of surplus among the players, when taking into account the worth of each coalition. To do this, he restricts attention to single-valued solutions and resorts to the axiomatic method. He proposes the following axioms on a single-valued solution:

1. Efficiency: The payoffs must add up to $v(N)$, which means that all the grand coalition surplus is allocated.
2. Symmetry: If two players are substitutes because they contribute the same to each coalition, the solution should treat them equally.
3. Additivity: The solution to the sum of two TU games must be the sum of what it awards to each of the two games.
4. Dummy player: If a player contributes nothing to every coalition, the solution should pay him nothing.

(To be precise, the name of the first axiom should be different. In an economic sense, the statement does imply efficiency in superadditive games, i.e., when for every pair of disjoint coalitions S and T , $v(S) + v(T) \leq v(S \cup T)$. In the absence of superadditivity, though, forming the grand coalition is not necessarily efficient, because a higher aggregate payoff can be obtained from a different coalition structure.)

The surprising result in Shapley (1953) is this:

Theorem 5 (Shapley (1953)) There is a unique single-valued solution to TU games satisfying efficiency, symmetry, additivity and dummy. It is what today we call the Shapley value, the function that assigns to each player i the payoff

$$\text{Sh}_i(N, v) = \sum_{S, i \in S} \frac{(|S| - 1)!(|N| - |S|)!}{|N|!} [v(S) - v(S \setminus \{i\})].$$

That is, the Shapley value awards to each player the average of his marginal contributions to each coalition. In taking this average, all orders of the players are considered to be equally likely. Let us assume, also without loss of generality, that $v(\{i\}) = 0$ for each player i .

What is especially surprising in Shapley's result is that nothing in the axioms (with the possible exception of the dummy axiom) hints at the idea of marginal contributions, so marginality in general is the outcome of all the axioms, including additivity or linearity. Among the axioms utilized by Shapley, additivity is the one with a lower normative content: it is simply a mathematical property to justify simplicity in the computation of the solution. Young in (Young 1985) provides a beautiful counterpart to Shapley's theorem. He drops additivity (as well as the dummy player axiom), and instead, uses an axiom of marginality. Marginality means that the solution should pay the same to a player in two games if his or her marginal contributions to coalitions is the same in both games. Marginality is an idea with a strong tradition in economic theory. Young's result is "dual" to Shapley's, in the sense that marginality is assumed and additivity derived as the result:

Theorem 6 (Young (1985)) There exists a unique single-valued solution to TU games satisfying efficiency, symmetry and marginality. It is the Shapley value.

Apart from these two, (Hart and Mas-Colell 1989) provides further axiomatizations of the Shapley value using the idea of potential and the concept of consistency, as described in the previous section.

There is no single way to extend the Shapley value to the class of NTU games. There are three

main extensions that have been proposed: the Shapley λ -transfer value (Shapley 1969), the Harsanyi value (Harsanyi 1963), and the Maschler-Owen consistent value (Maschler and Owen 1992). They were axiomatized in Aumann (1985), de Clippel et al. (2004), Hart (1985), respectively.

The Connections with Competitive Equilibrium

As was the case for the core, there is a value equivalence theorem. The result holds for the TU domain (see Aumann 1975; Aumann and Shapley 1974; Shapley 1964). It can be shown that the Shapley value payoffs can be supported by competitive prices. Furthermore, in large enough economies, the set of competitive payoffs “shrinks” to approximate the Shapley value. However, the result cannot be easily extended to the NTU domain. While it holds for the λ -transfer value, it need not obtain for the other extensions. For further details, the interested reader is referred to Hart (2008) and the references therein.

Non-cooperative Implementation

Reference Gul (1989) was the first to propose a procedure that provided some non-cooperative foundations of the Shapley value. Later, other authors have provided alternative procedures and techniques to the same end, including (Hart and Mas-Colell 1996; Krishna and Serrano 1995; Pérez-Castrillo and Wettstein 2001; Winter 1994).

We shall concentrate on the description of the procedure proposed by Hart and Mas-Colell in Hart and Mas-Colell (1996). Generalizing an idea found in Mas-Colell (1988), which studies the case of $\delta = 0$ – see below –, Hart and Mas-Colell propose the following non-cooperative procedure. With equal probability, each player $i \in N$ is chosen to publicly make a feasible proposal to the others: $(x_1, \dots, x_n)$ is such that the sum of its components cannot exceed $v(N)$. The other players get to respond to it in sequence, following a prespecified order. If all accept, the proposal is implemented; otherwise, a random device is triggered. With probability $0 \leq \delta < 1$ the same game continues being played among the same n players (and thus, a new proposer will be

chosen again at random among them), but with probability $1 - \delta$, the proposer leaves the game. He is paid 0 and his resources are removed, so that in the next period, proposals to the remaining $n - 1$ players cannot add up to more than $v(N \setminus \{i\})$. A new proposer is chosen at random among the set $N \setminus \{i\}$, and so on.

As shown in Hart and Mas-Colell (1996), there exists a unique stationary self-enforcing prediction of this procedure, and it actually coincides with the Shapley value payoffs for any value of δ . (Stationarity means that strategies cannot be history dependent). As $\delta \rightarrow 1$, the Shapley value payoffs are also obtained not only in expectation, but with independence of who is the proposer. One way to understand this result, as done in Hart and Mas-Colell (1996), is to check that the rules of the procedure and stationary behavior in it are in agreement with Shapley’s axioms. That is, the equilibrium relies on immediate acceptances of proposals, stationary strategies treat substitute players similarly, the equations describing the equilibrium have an additive structure, and dummy players will have to receive 0 because no resources are destroyed if they are asked to leave. It is also worth stressing the important role in the procedure of players’ marginal contributions to coalitions: following a rejection, a proposer incurs the risk of being thrown out and the others of losing his resources, which seem to suggest a “price” for them.

In Krishna and Serrano (1995), the authors study the conditions under which stationarity can be removed to obtain the result. Also, Pérez-Castrillo and Wettstein (2001) uses a variant of the Hart and Mas-Colell procedure, by replacing the random choice of proposers with a bidding stage, in which players bid to obtain the right to make proposals.

An Application

Consider again the class of majority problems in a parliament consisting of 100 seats. As we shall see, the Shapley value is a good way to understand the power that each party has in the legislature.

Let us begin by considering again the problem in which player 1 has 90 seats, while each of the five small parties has 2 seats. It is easy to see that

the Shapley value, like the core in this case, awards the entire unit of surplus to player 1: effectively, each of the small parties is a dummy player, and hence, the Shapley value awards zero to each of them.

Consider a second problem, in which player 1 is a big party with 35 seats, and there are 5 small parties, with 13 seats each. The Shapley value awards $1/3$ to the large party, and, by symmetry, $2/15$ to each of the small parties. To see this, we need to see when the marginal contributions of player 1 to any coalition are positive. Recall that there are $6!$ possible orders of players. Note how, if player 1 arrives first or second in the room in which the coalition is forming, his marginal contribution is zero: the coalition was losing before he arrived and continues to be a losing coalition after his arrival. Similarly, his marginal contribution is also zero if he arrives fifth or sixth to the coalition; indeed, in this case, before he arrives the coalition is already winning, so he adds nothing to it. Thus, only when he arrives third or fourth, which happens a third of the times, does he change the nature of the coalition, from losing to winning. This explains his Shapley value share of $1/3$. In this game, the Shapley value payoffs roughly correspond to the proportion of seats that each party has.

Next, consider a third problem in which there are two large parties, while the other four parties are very small. For example, let each of the large parties have 48 seats (say, players 1 and 2), while each of the four small parties has only one seat. Now, the Shapley value payoffs are 0.3 to each of the two large parties, and 0.1 to each of the small ones. To see this, note that the marginal contribution of a small party is only positive when he comes fourth in line, and out of the preceding three parties in the coalition, exactly one of them is a large party, i.e., 72 orders out of the $5!$ orders in which he is fourth. That is, $(72/5!) \times (1/6) = 1/10$. In this case, the competition between the large parties for the votes of the small parties increases the power of the latter quite significantly, with respect to the proportion of seats that each of them holds.

Finally, consider a fourth problem with two large parties (players 1 and 2) with 46 seats

each, one mid-size party (player 3) with 5 seats, and three small parties, each with one seat. First, note that each of the three small parties has become a dummy player: no winning coalition where he belongs becomes losing if he leaves the coalition, and so players 4, 5 and 6 are paid zero by the Shapley value. Now, note that, despite the substantial difference of seats between each large party and the mid-size party, each of them is identical in terms of marginal contributions to a winning coalition. Indeed, for $i = 1, 2, 3$ player i 's marginal contribution to a coalition is positive only if he arrives second or third or fourth or fifth (and out of the preceding players in the coalition, exactly one is one of the non-dummy players). Note how the Shapley value captures nicely the changes in the allocation of power due to each different political scenario. In this case, the fierce competition between the large parties for the votes of player 3, the swinging party to form a majority, explains the equal share of power among the three.

Future Directions

This article has been a first approach to cooperative game theory, and has emphasized two of its most important solution concepts. The literature on these topics is vast, and the interested reader is encouraged to consult the general references listed below. For the future, one should expect to see progress of the theory into areas that have been less explored, including games with asymmetric information and games with coalitional externalities. In both cases, the characteristic function model must be enriched to take care of the added complexities.

Relevant to this encyclopedia are issues of complexity. The complexity of cooperative solution concepts has been studied (see, for instance, Deng and Papadimitriou 1994). In terms of computational complexity, the Shapley value seems to be easy to compute, while the core is harder, although some classes of games have been identified in which this task is also simple.

Finally, one should insist on the importance of novel and fruitful applications of the theory to

shed new light on concrete problems. In the case of the core, for example, the insights of core stability in matching markets have been successfully applied by Alvin Roth and his collaborators to the design of matching markets in the “real world” (e.g., the job market for medical interns and hospitals, the allocation of organs from donors to patients, and so on) – see Roth (2002).

Bibliography

Primary Literature

- Anderson RM (1978) An elementary core equivalence theorem. *Econometrica* 46:1483–1487
- Anderson RM (2008) Core convergence. In: Durlauf S, Blume L (eds) *The new Palgrave dictionary of economics*, 2nd edn. McMillan, London
- Aumann RJ (1964) Markets with a continuum of traders. *Econometrica* 32:39–50
- Aumann RJ (1975) Values of markets with a continuum of traders. *Econometrica* 43:611–646
- Aumann RJ (1985) An axiomatization of the non-transferable utility value. *Econometrica* 53:599–612
- Aumann RJ (1987) Game theory. In: Eatwell J, Milgate M, Newman P (eds) *The new Palgrave dictionary of economics*. Norton, New York
- Aumann RJ, Peleg B (1960) Von Neumann-Morgenstern solutions to cooperative games without side payments. *Bull Am Math Soc* 66:173–179
- Aumann RJ, Shapley LS (1974) *Values of non-atomic games*. Princeton University Press, Princeton
- Bondareva ON (1963) Some applications of linear programming methods to the theory of cooperative games (in Russian). *Problemy Kibernetiki* 10(119):139
- de Clippel G, Peters H, Zank H (2004) Axiomatizing the Harsanyi solution, the symmetric egalitarian solution and the consistent solution for NTU-games. *I J Game Theory* 33:145–158
- Debreu G, Scarf H (1963) A limit theorem on the core of an economy. *Int Econ Rev* 4:235–246
- Deng X, Papadimitriou CH (1994) On the complexity of cooperative solution concepts. *Math Oper Res* 19:257–266
- Edgeworth FY (1881) *Mathematical Psychics*. Kegan Paul Publishers, London. (reprinted in 2003) Newman P (ed) *F. Y. Edgeworths Mathematical Psychics and Further Papers on Political Economy*. Oxford University Press, Oxford
- Gillies DB (1959) Solutions to general non-zero-sum games. In: Tucker AW, Luce RD (eds) *Contributions to the theory of games IV*. Princeton University Press, Princeton, pp 47–85
- Gul F (1989) Bargaining foundations of Shapley value. *Econometrica* 57:81–95
- Harsanyi JC (1963) A simplified bargaining model for the n -person cooperative game. *Int Econ Rev* 4:194–220
- Hart S (1985) An axiomatization of Harsanyi's non-transferable utility solution. *Econometrica* 53:1295–1314
- Hart S (2008) Shapley value. In: Durlauf S, Blume L (eds) *The new Palgrave dictionary of economics*, 2nd edn. McMillan, London
- Hart S, Mas-Colell A (1989) Potencial, value and consistency. *Econometrica* 57:589–614
- Hart S, Mas-Colell A (1996) Bargaining and value. *Econometrica* 64:357–380
- Krishna V, Serrano R (1995) Perfect equilibria of a model of n -person non-cooperative bargaining. *I J Game Theory* 24:259–272
- Mas-Colell A (1988) Algunos comentarios sobre la teoría cooperativa de los juegos. *Cuadernos Economicos* 40:143–161
- Maschler M, Owen G (1992) The consistent Shapley value for games without side payments. In: Selten R (ed) *Rational interaction: essays in honor of John Harsanyi*. Springer, New York
- Moldovanu B, Winter E (1995) Order independent equilibria. *Games Econ Behav* 9:21–34
- Nash JF (1953) Two person cooperative games. *Econometrica* 21:128–140
- Peleg B (1985) An axiomatization of the core of cooperative games without side payments. *J Math Econ* 14:203–214
- Peleg B (1986) On the reduced game property and its converse. *I J Game Theory* 15:187–200
- Pérez-Castrillo D (1994) Cooperative outcomes through non-cooperative games. *Games Econ Behav* 7:428–440
- Pérez-Castrillo D, Wettstein D (2001) Bidding for the surplus: a non-cooperative approach to the Shapley value. *J Econ Theory* 100:274–294
- Perry M, Reny P (1994) A non-cooperative view of coalition formation and the core. *Econometrica* 62:795–817
- Roth AE (2002) The economist as engineer: game theory, experimentation and computation as tools for design economics. *Econometrica* 70:1341–1378
- Scarf H (1967) The core of an N person game. *Econometrica* 38:50–69
- Serrano R (1995) A market to implement the core. *J Econ Theory* 67:285–294
- Serrano R (2005) Fifty years of the Nash program, 1953–2003. *Investigaciones Económicas* 29:219–258
- Serrano R, Vohra R (1997) Non-cooperative implementation of the core. *Soc Choice Welf* 14:513–525
- Serrano R, Volij O (1998) Axiomatizations of neoclassical concepts for economies. *J Math Econ* 30:87–108
- Shapley LS (1953) A value for n -person games. In: Tucker AW, Luce RD (eds) *Contributions to the theory of games II*. Princeton University Press, Princeton, pp 307–317
- Shapley LS (1964) Values of large games VII: a general exchange economy with money. Research Memorandum 4248-PR RAND Corporation, Santa Monica

- Shapley LS (1967) On balanced sets and cores. *Nav Res Logist Q* 14:453–460
- Shapley LS (1969) Utility comparison and the theory of games. In: *La Décision: Agrégation et Dynamique des Ordres de Préférence*. CNRS, Paris
- Shapley LS (1971) Cores of convex games. *I J Game Theory* 1:11–26
- von Neumann J, Morgenstern O (1944) *Theory of games and economic behavior*. Princeton University Press, Princeton
- Walras L (1874) *Elements of pure economics, or the theory of social wealth*. English edition: Jaffé W (ed) Reprinted in 1984 by Orion Editions, Philadelphia
- Winter E (1994) The demand commitment bargaining and snowballing of cooperation. *Econ Theory* 4:255–273
- Young HP (1985) Monotonic solutions of cooperative games. *I J Game Theory* 14:65–72

Books and Reviews

- Myerson RB (1991) *Game theory: an analysis of conflict*. Harvard University Press, Cambridge
- Osborne MJ, Rubinstein A (1994) *A course in game theory*. MIT Press, Cambridge
- Peleg B, Sudholter P (2003) *Introduction to the theory of cooperative games*, 2nd edn. Kluwer/Springer, Amsterdam/Berlin
- Roth AE, Sotomayor M (1990) *Two-sided matching: a study in game-theoretic modeling and analysis*. Cambridge University Press, Cambridge

Dynamic Games with an Application to Climate Change Models

Prajit K. Dutta

Department of Economics, Columbia University,
New York, NY, USA

Article Outline

Glossary

Definition of the Subject

Introduction

The Dynamic – or Stochastic – Game Model

Equilibrium

The Dynamic – or Stochastic – Game: Results

Existence

Characterization

Feasible Payoffs

Individually Rational Payoffs

Dynamics

Global Climate Change – Issues, Models

Models

Global Climate Change – Results

Global Pareto Optima

A Markov-Perfect Equilibrium: “Business as Usual”

All SPE

Generalizations

Future Directions

Bibliography

Glossary

Players The agents who take actions. These actions can be – depending on application – the choice of capital stock, greenhouse emissions, level of savings, level of Research & Development expenditures, price level, quality and quantity of effort, etc.

Strategies Full contingent plans for the actions that players take. Each strategy incorporates a

choice of action not just once but rather a choice of action for every possible decision node for the player concerned.

Payoffs The utility or returns to a player from playing a game. These payoffs typically depend on the strategies chosen – and the consequent actions taken – by the player herself as well as those chosen by the other players in the game.

Game horizon The length of time over which the game is played, i. e., over which the players take actions. The horizon may be finite – if there are only a finite number of opportunities for decision-making – or infinite – when there are an infinite number of decision-making opportunities.

Equilibrium A vector of strategies, one for each player in the game, such that no player can unilaterally improve her payoffs by altering her strategy, if the others’ strategies are kept fixed.

Climate change The consequence to the earth’s atmosphere of economic activities such as the production and consumption of energy that result in a build-up of greenhouse gases such as carbon dioxide.

Definition of the Subject

The study of dynamic games is an important topic within game theory. Dynamic games involve the study of problems that are (a) inherently dynamic in nature (even without a game-theoretic angle) and (b) are naturally studied from a strategic perspective. Towards that end the structure generalizes dynamic programming – which is the most popular model within which inherently dynamic but non-strategic problems are studied. It also generalizes the model of repeated games within which strategic interaction is often studied but which structure cannot handle dynamic problems. A large number of economic problems fit these two requirements.

In this paper we examine the dynamic game model. The structure is discussed in detail as well as its principal results. Then the paper introduces a leading important application, the economics of climate change. It is shown that the problem is best studied as a dynamic commons game. Some recent models and associated results are then discussed.

We begin the analysis with a recall of the familiar model of repeated games (whose main results have been presented elsewhere in this volume). That is followed by the generalization of that framework to the model of dynamic – also known as stochastic or Markovian – games. These games may be thought of as “repeated games with a state variable”. The presence of a state variable allows the analysis of situations where there is a feedback: 100% intrinsic to the problem, a situation – or “state” – that changes over time often on account of the players’ past actions. (In contrast to repeated games where an identical stage game is played period after period). Such a state variable maybe capital stock, level of technology, national or individual wealth or even environmental variables such as the size of natural resources or the stock of greenhouse gases. To provide a concrete illustration of the dynamic game concepts and results, this paper will provide a fairly detailed overview of ongoing research by a number of authors on the very current and important topic of the economics of global climate change.

Section “[The Dynamic – or Stochastic – Game Model](#)” recalls the repeated games structure, introduces the subject of dynamic games and presents the dynamic games model. Section “[The Dynamic – or Stochastic – Game: Results](#)” presents – mostly with proofs – the main results from the theory of dynamic games. Section “[Global Climate Change – Issues, Models](#)” then introduces the problem of climate change, argues why the dynamic game framework is appropriate for studying the problem and presents a family of models that have been recently studied by Dutta and Radner – and in a variant by Dockner, Long and Sorger. Finally, section “[Global Climate Change – Results](#)” presents the main results of these analyzes of the climate change problem. Future directions for research are discussed in section “[Future Directions](#)” while references are collected in section “Bibliography”.

Introduction

In this paper we examine the dynamic game model. The structure is discussed in detail as well as its principal results. Then the paper introduces a leading important application, the economics of climate change. It is shown that the problem is best studied as a dynamic commons game. Some recent models and associated results are then discussed.

The Dynamic – or Stochastic – Game Model

The most familiar model of dynamic interaction is the Repeated Game model (described elsewhere in this volume). In that set-up players interact every period for many periods – finite or infinite in number. At each period they play exactly the same game, i. e., they pick from exactly the same set of actions and the payoff consequence of any given action vector is identical. Put differently, it is as if there is an intrinsic static set-up, a “state” of the system that never changes. The only thing that changes over time is (potentially) every player’s response to that fixed state, i. e., players (can) treat the game dynamically if they so wish but there is no inherent non-strategic reason to do so. An impartial “referee” choosing on behalf of the players to achieve some optimization aim indeed would pick the same action every period.

Things change in a set-up where the state can change over time. That is the structure to which we now turn. This set-up was introduced by Shapley (1953) under the name of Stochastic Games. It has since also been called Markovian Games – on account of the Markovian structure of the intrinsic problem – or Dynamic Games. We will refer to the set-up (for the most part) as Dynamic Games.

Set-Up

There are I players, and time is discrete. Each period the players interact by picking an action. Their action interaction take place at a given state which state changes as a consequence of the action interaction. There is a payoff that each

player receives in each period based on the action vector that was picked and the state.

The basic variables are:

Definition

t Time period $(0, 1, 2, \dots, T)$.

i Players $(1, \dots, I)$.

$s(t)$ State at the beginning of period t , $s(t) \in S$.

$a_i(t)$ Action taken by player i in period, $a_i(t) \in A_i$.

$a(t)$ $(a_1(t), a_2(t), \dots, a_I(t))$ Vector actions taken in period t .

$\pi_i(t)$ $\pi_i(s(t), a(t))$ Payoff of player i in period t .

$q(t)$ $q(s(t+1)|s(t), a(t))$ Conditional distribution of state at the beginning of period $t+1$.

δ The discount factor, $\delta \in [0, 1)$.

The state variable affects play in two ways as stated above. In any given period, the payoff to a player depends not only on the actions that she and other players take but it also depends on the state in that period. Furthermore, the state casts a shadow on future payoffs in that it evolves in a Markovian fashion with the state in the next period being determined – possibly stochastically – by the state in the current period and the action vector played currently.

The initial value of the state, $s(0)$, is exogenous. So is the discount factor δ and the game horizon, T . Note that the horizon can be finite or infinite. All the rest of the variables are endogenous, with each player controlling its own endogenous variable, the actions. Needless to add, both state as well as action variables can be multi-dimensional and when we turn to the climate change application it will be seen to be multi-dimensional in natural ways.

Example 1 S infinite – The state space can be countably or uncountably infinite. It will be seen that the infinite case, especially the uncountably infinite one, has embedded within it a number of technical complications and – partly as a consequence – much less is known about this case.

S finite – In this case, imagine that we have a repeated game like situation except that there are a

finite number of stage games any one of which gets played at a time.

Example 2 When the number of players is one, i. e., $I = 1$, then we have a dynamic programming problem. When the number of states is one, i. e., $\#(S) = 1$, then we have a repeated game problem. (Alternatively, repeated games constitute the special case where the conditional distribution brings a state s always back to itself, regardless of action). Hence these two very familiar models are embedded within the framework of dynamic games.

Histories and Strategies

Preliminaries – A *history* at time t , $h(t)$, is a list of prior states and action vectors up to time t (but not including $a(t)$)

$$h(t) = s(0), a(0), s(1), a(1), \dots, s(t).$$

Let the set of histories be denoted $H(t)$. A *strategy* for player i at time t , $\sigma_i(t)$, is a complete conditional plan that specifies a choice of action for every history. The choice may be probabilistic, i. e., may be an element of $P(A_i)$, the set of distributions over A_i . So a strategy at time t is

$$\sigma_i(t) : H(t) \rightarrow P(A_i).$$

A strategy for the entire game for player i , σ_i , is a list of strategies, one for every period: $\sigma_i = (\sigma_i(0), \sigma_i(1), \dots, \sigma_i(t), \dots)$. Let $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_I)$ denote a vector of strategies, one for each player.

A particular example of a strategy for player i is a pure strategy σ_i where $\sigma_i(t)$ is a deterministic choice (from A_i). This choice may, of course, be conditional on history, i. e., may be a map from $H(t)$ to A_i . Another example of a strategy for player i is one where the player's choice $\sigma_i(t)$ may be probabilistic but the conditioning variables are not the entire history but rather only the current state. In other words such a strategy is described by a map from S to $P(A_i)$ – and is called a Markovian strategy. Additionally, when the map is independent of time, the strategy is called a stationary Markovian strategy, i. e., a stationary

Markovian strategy for player i is described by a mapping: $f_i : S \rightarrow P(A_i)$.

Example 3 Consider, for starters, a pure strategy vector σ , i. e., a pure strategy choice for every i . Suppose further that q the conditional distribution on states is also deterministic. In that case, there is, in a natural way, a unique history that is generated by σ :

$$h(t; \sigma) = s(0), \quad a(0; \sigma), \quad s(1; \sigma), \\ a(1; \sigma), \dots, s(t; \sigma)$$

where $a(\tau; \sigma) = \sigma(\tau; h(\tau; \sigma))$ and $s(\tau + 1; \sigma) = q(s(\tau + 1) | s(\tau; \sigma), a(\tau; \sigma))$. This unique history associated with the strategy vector σ is also called the *outcome path* for that strategy. To every such outcome path there is an associated lifetime payoff

$$R_i(\sigma) = \sum_{t=0}^T \delta^t \pi_i(s(t; \sigma), a(t; \sigma)). \quad (1)$$

If σ is a mixed strategy, or if the conditional distribution q , is not deterministic, then there will be a joint distribution on the set of histories $H(t)$ generated by the strategy vector σ and the conditional distribution q in the obvious way. Moreover, there will be a marginal distribution on the state and action in period t , and under that marginal, an expected payoff $\pi_i(s(t; \sigma), a(t; \sigma))$. Thereafter lifetime payoffs can be written exactly as in Eq. 1. Consider the game that remains after every history $h(t)$. This remainder is called a *subgame*. The restriction of the strategy vector σ to the subgame that starts after history $h(t)$, is denoted $\sigma | h(t)$.

Equilibrium

A strategy vector σ^* is said to be a Nash Equilibrium (or NE) of the game if

$$R_i(\sigma^*) \geq R_i(\sigma_i, \sigma_{-i}^*), \quad \text{for all } i, \sigma_i. \quad (2)$$

A strategy vector σ^* is said to be a Subgame Perfect (Nash) Equilibrium of the game – referred

to in short as SPE – if not only is Eq. 2 true for σ^* but it is true for every restriction of the strategy vector σ^* to every subgame $h(t)$, i. e., is true for $\sigma^* | h(t)$ as well. In other words, σ^* is a SPE if

$$R_i(\sigma^* | h(t)) \geq R_i(\sigma_i, \sigma_{-i}^* | h(t)), \quad \text{for all } i, \sigma_i, h(t). \quad (3)$$

As is well-known, not all NE satisfy the further requirement of being a SPE. This is because a NE only considers the outcome path associated with that strategy vector σ^* – or, when the outcome path is probabilistic, only considers those outcome paths that have a positive probability of occurrence. That follows from the inequality Eq. 2. However, that does not preclude the possibility that players may have no incentive to follow through with σ^* if some zero probability history associated with that strategy is reached. (Such a history may be reached either by accident or because of deviation/experimentation by some player). In turn that may have material relevance because how players behave when such a history is reached will have significance for whether or not a player wishes to deviate against σ^* . Equation 3 ensures that – even after a deviation – σ^* will get played and that deviations are unprofitable.

Recall the definition of a stationary Markovian strategy (SMS) above. Associated with that class of strategies is the following definition of equilibrium. A stationary Markov strategy vector f^* is a Markov Perfect Equilibrium (MPE) if

$$R_i(f^*) \geq R_i(f_i, f_{-i}^*), \quad \text{for all } i, f_i.$$

Hence, a MPE restricts attention to SMS both on and off the outcome path. Furthermore, it only considers – implicitly – histories that have a positive probability of occurrence under f^* . Neither “restriction” is a restriction when T is infinite because when all other players play a SMS player i has a stationary dynamic programming problem to solve in finding his most profitable strategy and – as is well-known – he loses no payoff possibilities in restricting himself to SMS as well. And that best strategy is a best strategy on histories that have zero probabilities of occurrence

as well as histories that have a positive probability of occurrence. In particular therefore, when T is infinite, a MPE is also a SPE.

The Dynamic – or Stochastic – Game: Results

The main questions that we will now turn to are:

1. Is there always a SPE in a dynamic – or stochastic – game?
2. Is there a characterization for the set of SPE akin to the Bellman optimality equation of dynamic programming? If yes, what properties can be deduced of the SPE payoff set?
3. Is there a Folk Theorem for dynamic games – akin to that in Repeated Games?
4. What are the properties of SPE outcome paths?

The answers to questions 1–3 are very complete for finite dynamic games, i. e., games where the state space S is finite. The answer is also complete for questions 1 and 2 when S is countably infinite but when the state space is uncountably infinite, the question is substantively technically difficult and there is reason to believe that there may not always be a SPE. The finite game arguments for question 3 is conceptually applicable when S is (countably or uncountably) infinite provided some technical difficulties can be overcome. That and extending the first two answers to uncountably infinite S remain open questions at this point. Not a lot is known about Question 4.

Existence

The first result is due to Parthasarathy (1973) and applies to the infinite horizon model, i. e., where $T = \infty$. When T is finite, the result and arguments can be modified in a straightforward way as will be indicated in the remarks following the proof.

Theorem 1 Suppose that S is countable, A_i are finite sets, and the payoff functions π_i are bounded. Suppose furthermore that T is infinite. Then there is a MPE (and hence a SPE).

Proof The proof will be presented by way of a fixed point argument. The domain for the fixed point will be the set of stationary Markovian strategies:

$$M_j = \left\{ f_i : S \rightarrow P(A_i), \quad \text{s.t for all } s, \right. \\ \left. \sum_{a_i} f_i(a_i; s) = 1, f_i(a_i; s) \geq 0 \right\}.$$

Properties of M_j : In the pointwise convergence topology, M_i is compact. That this is so follows from a standard diagonalization argument by way of which a subsequence can be constructed from any sequence of SMS f_i^n such that the subsequence, call it $f_i^{n'}$ has the property that $f_i^{n'}(s)$ converges to some $f_i^0(s)$ for every s . Clearly, $f_i^0 \in M_i$. The diagonalization argument requires S to be countable and A_i to be finite.

M_i is also clearly convex since its elements are probability distributions on A_i at every state.

The mapping for which we shall seek a fixed point is the best response mapping:

$$B_i(f) = \{g_i \in M_i : R_i(g_i, f_{-i}) \geq R_i(f_i, f_{-i}), \quad \text{for all } f_i\}.$$

Since the best response problem for player i is a stationary dynamic programming problem, it follows that there is an associated value function for the problem, say v_i , such that it solves the optimality equation of dynamic programming

$$v_i(s) = \max_{\lambda_i} \left\{ \pi_i(s, \lambda_i, f_{-i}(s)) \right. \\ \left. + \delta \sum_{s'} v_i(s') q(s' | s, \lambda_i, f_{-i}(s)) \right\} \quad (4)$$

where

$$\pi_i(s, \lambda_i, f_{-i}(s)) = \sum_{a_i} \left[\sum_{a_{-i}} \pi_i(s, a_i, a_{-i}) \lambda_i(a_i) \right] \\ f_i(a_{-i}, s) \quad (5)$$

where $\lambda_i(a_i)$ is the probability of player i picking action a_i whilst $f_i(a_{-i}, s)$ is the product

probability of players other than i picking the action vector $\mathbf{a}_{-i}$. Similarly,

$$q(s' | s, \lambda_i, f_{-i}(s)) = \sum_{\mathbf{a}_{-i}} \left[\sum_{\mathbf{a}_i} q(s' | s, \mathbf{a}_i, \mathbf{a}_{-i}) \lambda_i(\mathbf{a}_i) \right] f_i(\mathbf{a}_{-i}, s). \quad (6)$$

Additionally, it follows that the best best response, i. e., g_i , solves the optimality equation, i. e.,

$$v_i(s) = \pi_i(s, g_i, f_{-i}(s)) + \delta \sum_{s'} v_i(s') q(s' | s, g_i, f_{-i}(s)) \quad (7)$$

where $\pi_i(s, g_i, f_{-i}(s))$ and $q(s' | s, g_i, f_{-i}(s))$ have the same interpretations as given by Eqs. 5 and 6.

Properties of B_i : B_i is a correspondence that is convex-valued and upper hemi-continuous. That B_i is convex-valued follows from the fact that we are operating in the set of mixed strategies, that every convex combination of mixed strategies is itself a mixed strategy and that every convex combination of best responses is also a best response.

To show that B_i is upper hemi-continuous, consider a sequence of other players' strategies f_{-i}^n , an associated best response sequence of player i , g_i^n with value function sequence v_i^n . Note that each of these best responses g_i^n satisfies the Eqs. 6 and 7 (for the value function v_i^n). By diagonalization there exist subsequences and subsequential pointwise convergent limits: $f_{-i}^n \rightarrow f_{-i}^0$, $g_i^n \rightarrow g_i^0$, and $v_i^n \rightarrow v_i^0$. It suffices to show that

$$v_i^0(s) = \max_{\lambda_i} \left\{ \pi_i(s, \lambda_i, f_{-i}^0(s)) + \delta \sum_{s'} v_i^0(s') q(s' | s, \lambda_i, f_{-i}^0(s)) \right\} \quad (8)$$

and

$$v_i^0(s) = \pi_i(s, g_i^0, f_{-i}^0(s)) + \delta \sum_{s'} v_i^0(s') q(s' | s, g_i^0, f_{-i}^0(s)). \quad (9)$$

Equation 9 will be proved by using the analog of Eq. 7, i. e.,

$$v_i^n(s) = \pi_i(s, g_i^n, f_{-i}^n(s)) + \delta \sum_{s'} v_i^n(s') q(s' | s, g_i^n, f_{-i}^n(s)). \quad (10)$$

Clearly the left-hand side of Eq. 10 converges to the left-hand side of Eq. 9. Lets check the right-hand side of each equation. Evidently

$$\begin{aligned} \sum_{\mathbf{a}_{-i}} \left[\sum_{\mathbf{a}_i} \pi_i(s, \mathbf{a}_i, \mathbf{a}_{-i}) g_i^n(\mathbf{a}_i) \right] f_i^n(\mathbf{a}_{-i}; s) \\ \rightarrow \sum_{\mathbf{a}_{-i}} \left[\sum_{\mathbf{a}_i} \pi_i(s, \mathbf{a}_i, \mathbf{a}_{-i}) g_i^0(\mathbf{a}_i) \right] f_i^0(\mathbf{a}_{-i}; s) \end{aligned}$$

since each component of the sum converges and we have a finite sum. Finally,

$$\begin{aligned} & \left| \sum_{s'} v_i^n(s') q(s' | s, g_i^n, f_{-i}^n(s)) - \sum_{s'} v_i^0(s') q(s' | s, g_i^0, f_{-i}^0(s)) \right| \\ & \leq \left| \sum_{s'} [v_i^n(s') - v_i^0(s')] q(s' | s, g_i^n, f_{-i}^n(s)) \right| \\ & \quad + \left| \sum_{s'} v_i^0(s') [q(s' | s, g_i^n, f_{-i}^n(s)) - q(s' | s, g_i^0, f_{-i}^0(s))] \right|. \end{aligned} \quad (11)$$

The first term in the right-hand side of the inequality above goes to zero by the dominated convergence theorem. The second term can be re-written as

$$\begin{aligned} \sum_{s'} v_i^0(s') q(s' | s, \mathbf{a}_i, \mathbf{a}_{-i}) \\ \cdot \sum_{\mathbf{a}_{-i}} \sum_{\mathbf{a}_i} [g_i^n(\mathbf{a}_i) f_i^n(\mathbf{a}_{-i}; s) - g_i^0(\mathbf{a}_i) f_i^0(\mathbf{a}_{-i}; s)] \end{aligned}$$

and goes to zero because each of the finite number of terms in the summation over action probabilities goes to zero. Hence the RHS of Eq. 10 converges to the RHS of Eq. 9 and the proposition is proved.

Remark 1 Note that the finiteness of A_i is crucial. Else, the very last argument would not go through, i. e., knowing that $[g_i^n(a_i)f_i^n(a_{-i}; s) - g_i^0(a_i)f_i^0(a_{-i}; s)] \rightarrow 0$ for every action vector $\mathbf{a}$ would not guarantee that the sum would converge to zero as well.

Remark 2 If the horizon were finite one could use the same argument to prove that there exists a Markovian strategy equilibrium, though not a stationary Markovian equilibrium. That proof would combine the arguments above with backward induction. In other words, one would first use the arguments above to show that there is an equilibrium at every state in the last period T . Then the value function so generated, v_i^T , would be used to show that there is an equilibrium in period $T - 1$ using the methods above thereby generating the relevant value function for the last two periods, v_i^{T-1} . And so on.

The natural question to ask at this point is whether the restriction of countable finiteness of S can be dropped (and – eventually – the finiteness restriction on A_i). The answer, unfortunately, is not easily. The problems are two-fold:

1. *Sequential Compactness of the Domain Problem* – If S is uncountably infinite, then it is difficult to find a domain M_i that is sequentially compact. In particular, diagonalization arguments do not work to extract candidate strategy and value function limits.
2. *Integration to the Limit Problem* – Note as the other players change their strategies, f_{-i}^n , continuation payoffs to player i change in two ways. They change first because the value function v_i^n changes, i. e., $v_i^n \neq v_i^m$ if $n \neq m$. Second, the expected continuation value changes because the measure over which the value function is being integrated, $q(s' | s, \lambda_i, f_{-i}^n(s))$, itself changes, i. e., $q(s' | s, \lambda_i, f_{-i}^n(s)) \neq q(s' | s, \lambda_i, f_{-i}^m(s))$. This is the well-known – and difficult – integration to the limit problem: simply knowing that v_i^n “converges” to v_i^0 in some sense – such as pointwise – and knowing that the integrating measure q^n “converges” to q^0 in some sense – such as in the weak topology – does not, in general, imply that

$$\int v_i^n dq^n \rightarrow \int v_i^0 dq^0. \quad (12)$$

(Of course in the previous sentence q^n is a more compact stand-in for $q(s' | s, \lambda_i, f_{-i}^n(s))$ and q^0 for $q(s' | s, \lambda_i, f_{-i}^0(s))$). There are a limited number of cases where Eq. 12 is known to be true. These results typically require q^n to converge to q^0 in some strong sense. In the dynamic game context what this means is that very strong convergence restrictions need to be placed on the transition probability q . This is the underlying logic behind results reported in (Duffie et al. 1994; Mertens and Parthasarathy 1987; Nowak 1985; Rieder 1979).

Such strong convergence properties are typically not satisfied when q is deterministic – which case comprises the bulk of the applications of the theory. Indeed simply imposing continuity when q is deterministic appears not to be enough to generate an existence result. Harris et al. (1995) and Dutta and Sundaram (1993) contain results that show that there may not be a SPE in finite horizon dynamic games when the transition function q is continuous. Whether other often used properties of q and π_i – such as concavity and monotonicity – can be used to rescue the issue remains an open question.

Characterization

The Bellman optimality equation has become a workhorse for dynamic programming analysis. It is used to derive properties of the value function and the optimal strategies. Moreover it provides an attractive and conceptually simple way to view a multiple horizon problem as a series of one-stage programming problems by exploiting the recursive structure of the optimization set-up. A natural question to ask, since dynamic games are really multi-player versions of dynamic programming, is whether there is an analog of the Bellman equation for these games. Abreu et al. – APS (1990), in an important and influential paper, showed that this is indeed the case for repeated games. They defined an operator, hereafter the APS operator, whose largest fixed point is the set of SPE payoffs in a

repeated game and whose every fixed point is a subset of the set of SPE payoffs. (Thereby providing a necessary and sufficient condition for SPE equilibrium payoffs in much the same way that the unique fixed point of the Bellman operator constitutes the value function for a dynamic programming problem). As with the Bellman equation, the key idea is to reduce the multiple horizon problem to a (seemingly) static problem.

In going from repeated to dynamic games there are some technical issues that arise. We turn now to that analysis pointing out along the way where

the technical pitfalls are. Again we start with the infinite horizon model, i. e., where $T = \infty$. When T is finite, the result and arguments can be modified in a straightforward way as will be indicated in a remark following the proof.

But first, some definitions. Suppose for now that S is countable.

APS Operator – Consider a compact-valued correspondence, W defined on domain S which takes values that are subsets of R^I . Define the APS operator on W , call it LW , as follows:

$$LW(s) = \left\{ \begin{array}{l} v \in R^I : \exists \hat{f} \in P(A) \quad \text{and} \\ w : S \times A \times S \rightarrow W, \\ \text{uniformly bounded, s.t.} \quad v_i = \pi_i(s, \hat{f}) \\ + \delta \sum_{s'} w_i(s, \hat{f}, s') q(s' | s, \hat{f}) \\ \geq \pi_i(s, a_i, \hat{f}_{-i}) + \delta \sum_{s'} w_i(s, a_i, \hat{f}_{-i}, s') \\ q(s' | s, a_i, \hat{f}_{-i}), \quad \text{for all } a_i, i \end{array} \right\} \quad (13)$$

where, as before,

$$\pi_i(s, \hat{f}) = \sum_{a_{-i}} \left[\sum_{a_i} \pi_i(s, a_i, a_{-i}) \hat{f}_i(a_i) \right] \hat{f}_{-i}(a_{-i}; s) \quad (14)$$

and

$$q(s' | s, \hat{f}) = \sum_{a_{-i}} \left[\sum_{a_i} q(s' | s, a_i, a_{-i}) \hat{f}_i(a_i) \right] \hat{f}_{-i}(a_{-i}; s) \quad (15)$$

Finally, let V^* denote the SPE payoffs correspondence, i. e., $V^*(s)$ is the set of SPE payoffs starting from initial state s . Note that the correspondence is non-empty by virtue of Theorem 1.

Theorem 2 Suppose that S is countable, A_i are finite sets, and the payoff functions π_i are bounded. Suppose furthermore that T is infinite. Then (a) V^* is a fixed point of the APS operator, i. e., $LV^* = V^*$. Furthermore, (b) consider any other fixed point, i. e., a correspondence $\tilde{V}$ such that $L\tilde{V} = \tilde{V}$. Then it must be the case that $\tilde{V} \subset V^*$.

Finally, (c) there is an algorithm that generates the SPE correspondence, V^* .

Proof of a: Suppose that $v^* \in V^*(s)$, i. e., is a SPE payoff starting from s . Then, by definition, there is a first-period play, f^* and a continuation strategy after every one-period history $h(1)$, $\sigma^*(h(1))$, such that $\{f^*, \sigma^*(h(1))\}$ is a SPE. By definition, then, the payoffs associated with each history-dependent strategy, $\sigma^*(h(1))$, call them $w_i(s, f^*, s')$, satisfy Eq. 13. (Note that π_i bounded implies that the lifetime payoffs $w_i(s, f^*, s')$ are uniformly bounded). In other words, $v^* \in LV^*(s)$. On the other hand, suppose that $v^* \in LV^*(s)$. Then, by definition, there is a first-period play, f^* , and SPE payoffs, $w(s, f^*, s')$, that together satisfy Eq. 13. Let the SPE strategy associated with $w(s, f^*, s')$ be $\sigma^*(h(1))$. It is not difficult to see that the concatenated strategy – f^* , $\sigma^*(h(1))$ – forms a SPE. Since the associated lifetime payoff is v^* , it follows that $v^* \in V^*(s)$. $\square$

Proof of b: Suppose that $L\tilde{V} = \tilde{V}$. Let $v \in L\tilde{V}$. Then, from Eq. 13, there is $\hat{f}$, and SPE payoffs, $w(s, \hat{f}, s')$ that satisfy the equation. In turn, since $w(s, \hat{f}, s') \in \tilde{V}(s') = L\tilde{V}(s')$ there is an associated $f(s')$ and $w(s', f, s'')$ for which Eq. 13

holds. By repeated application of this idea, we can create a sequence of strategies for periods $t = 0, 1, 2, \dots - \tilde{f}, f(s'), \vec{f}(s, s'), \dots$ such that at each period Eq. 13 holds. Call the strategy so formed, ϕ . This strategy can then not be improved upon by a single-period deviation. A standard argument shows that if a strategy cannot be profitably deviated against in one period then it cannot be profitably deviated against even by deviations in multiple periods. (This idea of “unimprovability” is already present in dynamic programming). Within the context of repeated games, it was articulated by Abreu (1988).

Proof of c: Note two properties of the APS operator:

Lemma 1 LW is a compact-valued correspondence (whenever W is compact-valued).

Proof Consider Eq. 13. Suppose that $v^n \in LW(s)$ for all n , with associated $\hat{f}^n$ and w^n . By diagonalization, there exists a subsequence s.t. $v^n \rightarrow v^0, \hat{f}^n \rightarrow \hat{f}^0$ and $w^n \rightarrow w^0$. This arguments uses the countability of S and the finiteness of A_i . From Eq. 14 evidently $\pi_i(s, \hat{f}^n) \rightarrow \pi_i(s, \hat{f}^0)$ and similarly from Eq. 15 $\sum_{s'} w_i(s, \hat{f}^n, s') q(s'|s, \hat{f}^n)$ goes to $\sum_{s'} w_i(s, \hat{f}^0, s') q(s'|s, \hat{f}^0)$. Hence the inequality in Eq. 13 is preserved and $v^0 \in LW(s)$. $\square$

It is not difficult to see that – on account of the boundedness of π_i – if W has a uniformly bounded selection, then so does LW . Note that the operator is also monotone in the set-inclusion sense, i. e., if $W'(s) \subset W(s)$ for all s then $LW' \subset LW$.

The APS algorithm finds the set of SPE payoffs by starting from a particular starting point, an initial set $W^0(s)$ that is taken to be the set of all feasible payoffs from initial state s . (And hence the correspondence W^0 is so defined for every initial state). Then define, $W^1 = LW^0$. More generally, $W^{n+1} = LW^n$, $n \geq 0$. It follows that $W^1 \subset W^0$. This is because W^1 requires a payoff that is not only feasible but additionally satisfies the incentive inequality of Eq. 13 as well. From the monotone inclusion property above it then follows that, more generally, $W^{n+1} \subset W^n$, $n \geq 0$. Furthermore, $W^n(s)$ is a non-empty, compact set

for all n (and s). Hence, $W^\infty(s) = \bigcap_n W^n(s) = \lim_{n \rightarrow \infty} W^n(s)$ is non-empty and compact.

Let us now show that W^∞ is a fixed point of the APS operator, i. e., that $LW^\infty = W^\infty$.

Lemma 2 $LW^\infty = W^\infty$, or, equivalently, $L(\lim_{n \rightarrow \infty} W^n) = \lim_{n \rightarrow \infty} LW^n$.

Proof Clearly, by monotonicity, $L(\lim_{n \rightarrow \infty} W^n) \subset \lim_{n \rightarrow \infty} LW^n$. So consider a $v \in LW^\infty(s)$, for all n . By Eq. 13 there is at each n an associated first-period play f^n and a continuation payoff $w^n(s, f^n, s')$ such that the inequality is satisfied and

$$v_i = \pi_i(s, f^n) + \delta \sum_{s'} w_i^n(s, f^n, s') q(s'|s, f^n).$$

By the diagonalization argument, and using the countability of S , we can extract a (subsequential) limit $f^\infty = \lim_{n \rightarrow \infty} f^n$ and $w^\infty = \lim_{n \rightarrow \infty} w^n$. Clearly, $w^\infty \in W^\infty$. Since equalities and inequalities are maintained in the limit, equally clearly

$$v_i = \pi_i(s, f^\infty) + \delta \sum_{s'} w_i^\infty(s, f^\infty, s') q(s'|s, f^\infty)$$

and

$$v_i \geq \pi_i(s, a_i, f_{-i}^\infty) + \delta \sum_{s'} w_i^\infty(s, a_i, f_{-i}^\infty, s') q(s'|s, a_i, f_{-i}^\infty), \text{ for all } a_i, i$$

thereby proving that $v \in L(\lim_{n \rightarrow \infty} W^n(s))$. The lemma is proved. $\square$

Since the set of SPE payoffs, $V^*(s)$, is a subset of $W^0(s)$ – and $LV^*(s) = V^*(s)$ – it further follows $V^*(s) \subset W^\infty(s)$, for all s . From the previous lemma, and part b), it follows that $V^*(s) \supset W^\infty(s)$, for all s . Hence, $V^* = W^\infty$. Theorem is proved. $\square$

A few remarks are in order.

Remark 1 If the game horizon T is finite, there is an immediate modification of the above arguments. In the algorithm above, take W^0 to be the set of SPE payoffs in the one-period game (with payoffs π_i for player i).

Use the APS operator thereafter to define $W^{n+1} = LW^n$, $n \geq 0$. It is not too difficult to show that W^n is the set of SPE payoffs for a game that lasts $n+1$ periods (or has n remaining periods after the first one).

Remark 2 Of course an immediate corollary of the above theorem is that the set of SPE payoffs $V^*(s)$ is a compact set for every initial state s . Indeed one can go further and show that V^* is in fact an upper hemi-continuous correspondence. The arguments are very similar to those used above – plus the Maximum Theorem.

Remark 3 Another way to think of Theorem 2 is that it is also an existence theorem. Under the conditions outlined in the result, the SPE equilibrium set has been shown to be non-empty. Of course this is not a generalization of Theorem 1 since Theorem 2 does not assert the existence of a MPE.

Remark 4 When the state space A_i is infinite or the state space S is uncountably infinite we run into technical difficulties. The complications arise from not being able to take limits. Also, as in the discussion of the Integration to the Limit problem, integrals can fail to be continuous thereby rendering void some of the arguments used above.

Folk Theorem

The folk theorem for Repeated Games – Fudenberg and Maskin (1986) following up on earlier contributions – is very well-known and the most cited result of that theory. It proves that the necessary conditions for a payoff to be a SPE payoff – feasibility and individual rationality – are also (almost) sufficient provided the discount factor δ is close enough to 1. This is the result that has become the defining result of Repeated Games. For supporters, the result and its logic of proof are a compelling demonstration of the power of reciprocity, the power of longterm relationships in fostering cooperation through the lurking power of “punishments” when cooperation breaks down. It is considered equally important and significant that such long-term relationships and behaviors are sustained through implicit promises and threats which therefore do

not violate any legal prohibitions against explicit contracts that specify such behavior. For detractors, the “anything goes” implication of the Folk Theorem is a clear sign of its weakness – or the weakness of the SPE concept – in that it robs the theory of all predictive content. Moreover there is a criticism, not entirely correct, that the strategies required to sustain certain behaviors are so complex that no player in a “real-world” setting could be expected to implement them.

Be that as it may, the Folk Theorem question in the context of Dynamic Games then is: is it the case that feasibility and individual rationality are also (almost) enough to guarantee that a payoff is a SPE payoff at high enough δ ? Two sets of obstacles arise in settling this question. Both emanate from the same source, the fact that the state does not remain fixed in the play of the games, as it does in the case of Repeated Games. First, one has to think long and hard as to how one should define individual rationality. Relatedly, how does one track feasibility? In both cases, the problem is that what payoff is feasible and individually rational depends on the state and hence changes after every history $h(t)$. Moreover, it also changes with the discount factor δ . The second set of problems stems from the fact that a deviation play can unalterably change the future in a dynamic game – unlike a repeated game where the basic game environment is identical every period. Consequently one cannot immediately invoke the logic of repeated game folk theorems which basically work because any deviation has only short-term consequences while the punishment of the deviation is long-term. (And so if players are patient they will not deviate).

Despite all this, there are some positive results that are around. Of these, the most comprehensive is one due to Dutta (1995). To set the stage for that result, we need a few crucial preliminary results. For this sub-section we will assume that S is finite – in addition to A_i .

Feasible Payoffs

Role of Markovian Strategies – Let $F(s, \delta)$ denote the set of “average” feasible payoffs from initial state s and for discount factor δ . By that I mean

$$F(s, \delta) = \left\{ v \in \mathbb{R}^I : \exists \text{ strategy } \sigma \right. \\ \left. \text{s.t. } v = (1 - \delta) \sum_{t=0}^T \delta^t \pi_i(s(t; \sigma), a(t; \sigma)) \right\}.$$

Let $\Phi(s, \delta)$ denote the set of “average” feasible payoffs from initial state s and for discount factor δ that are generated by pure stationary Markovian strategies – PSMS. Recall that a SMS is given by a map f_i from S to the probability distributions over A_i , so that at state $s(t)$ player i chooses the mixed strategy $f_i(s(t))$. A pure SMS is one where the map f_i is from S to A_i . In other words,

$$\Phi(s, \delta) = \left\{ v \in \mathbb{R}^I : \exists \text{ PSMS } f \right. \\ \left. \text{s.t. } v = (1 - \delta) \sum_{t=0}^T \delta^t \pi_i(s(t; f), a(t; f)) \right\}.$$

Lemma 3 Any feasible payoff in a dynamic game can be generated by averaging over payoffs to stationary Markov strategies, i. e., $F(s, \delta) = \text{co}\Phi(s, \delta)$, for all (s, δ) .

Proof Note that $F(s, \delta) = \text{co}$ [extreme points $F(s, \delta)$]. In turn, all extreme points of $F(s, \delta)$ are generated by an optimization problem of the form: $\max_{\sigma} \sum_{i=1}^I \alpha_i v_i(s, \delta)$. That optimization problem is a dynamic programming problem. Standard results in dynamic programming show that the optimum is achieved by some stationary Markovian strategy.

Let $F(s)$ denote the set of feasible payoffs under the long-run average criterion. The next result will show that this is the set to which discounted average payoffs converge:

Lemma 4 $F(s, \delta) \rightarrow F(s)$, as $\delta \rightarrow 1$, for all s .

Proof Follows from the fact that (a) $F(s) = \text{co}\Phi(s)$ where $\Phi(s)$ is the set of feasible long-run average payoffs generated by stationary Markovian strategies, and (b) $\Phi(s, \delta) \rightarrow \Phi(s)$. Part (b) exploits the finiteness of S (and A_i). $\square$

The lemmas above simplify the answer to the question: What is a feasible payoff in a dynamic game? Note that they also afford a dimensional reduction in the complexity and number of strategies that one needs to keep track of to answer the question. Whilst there are an uncountably infinite number of strategies – even with finite S and A_i – including the many that condition on histories in arbitrarily complex ways – the lemmas establish that all we need to track are the finite number of PSMS. Furthermore, whilst payoffs do depend on δ , if the discount factor is high enough then the set of feasible payoffs is well-approximated by the set of feasible long-run average payoffs to PSMS.

One further preliminary step is required however. This has to do with the fact that while $v \in F(s, \delta)$ can be exactly reproduced by a period 0 average over PSMS payoffs, after that period continuation payoffs to the various component strategies may generate payoffs that could be arbitrarily distant from v . This, in turn, can be problematical since one would need to check for deviations at every one of these (very different) payoffs. The next lemma addresses this problem by showing that there is an averaging over the component PSMS that is ongoing, i. e., happens periodically and not just at period 0, but which, consequently, generates payoffs that after all histories stays arbitrarily close to v .

For any two PSMS f^1 and f^2 denote a *time-cycle strategy* as follows: for T^1 periods play proceeds along f^1 , then it moves for T^2 periods to f^2 . After the elapse of the $T^1 + T^2$ periods play comes back to f^1 for T^1 periods and f^2 for T^2 periods. And so on. Define $\lambda^1 = T^1/(T^1 + T^2)$. In the obvious way, denote a general time-cycle strategy to be one that cycles over any finite number of PSMS f^k where the proportion of time spent at strategy f^k is λ^k and allows the lengths of time to depend on the initial state at the beginning of the cycle.

Lemma 5 Pick any $v \in \cap_s F(s)$. Then for all $\varepsilon > 0$ there is a time cycle strategy such that its long-run average payoffs within ε of v after all histories.

Proof Suppose that $v = \sum_k \lambda^k(s) v^k(s)$ where $v^k(s)$ is the long-run average payoff is the k th PSMS

when the initial state is s . Ensure that T^k is chosen such that (a) the average payoff over those periods under that PSMS $- 1/T^k \sum_{t=0}^{T^k-1} \pi_i(s(t, f^k), a(t, f^k))$ is within ε of $v^k(s)$ for all s . And (b) that $T^k(s)/\sum_l T^l(s)$ is arbitrarily close to $\lambda^k(s)$ for all s . $\square$

Since $\Phi(s, \delta) \rightarrow \Phi(s)$ it further follows that the above result also holds under discounting:

Lemma 6 Pick any $v \in \cap_s F(s)$. Then for all $\varepsilon > 0$ there is a time cycle strategy and a discount cut-off $\delta(\varepsilon) < 1$ such that the discounted average payoffs to that strategy are within ε of v for all $\delta > \delta(\varepsilon)$ and after all histories.

Proof Follows from the fact that $(1 - \delta)/((1 - \delta^T) \sum_{t=0}^{T^k-1} \delta^t \pi_i(s(t, f^k), a(t, f^k)))$ goes to $\frac{1}{T^k} \sum_{t=0}^{T^k-1} \pi_i(s(t, f^k), a(t, f^k))$ as $\delta \rightarrow 1$. $\square$

Individually Rational Payoffs

Recall that a min-max payoff is a payoff level that a player can guarantee by playing a best response. In a Repeated Game that is defined at the level of the component stage game. Since there is no analog of that in a dynamic game, the min-max needs to be defined over the entire game – and hence is sensitive to initial state and discount factor:

$$m_i(s, \delta) = \min_{\sigma_{-i}} \max_{\sigma_i} R_i(\sigma | s, \delta).$$

Evidently, given (s, δ) , in a SPE it cannot be that player i gets a payoff $v_i(s, \delta)$ that is less than $m_i(s, \delta)$. Indeed that inequality must hold at all states for a strategy to be a SPE, i.e., for all $s(t)$ it must be the case that $v_i(s(t), \delta) \geq m_i(s(t), \delta)$. But, whilst necessary, even that might not be a sufficient condition for the strategy to be a SPE. The reason is that if player i can deviate and take the game to, say, s' at $t + 1$, rather than $s(t + 1)$, he would do so if $v_i(s(t), \delta) < m_i(s', \delta)$ since continuation payoffs from s' have to be at least as large as the latter level and this deviation would be worth essentially that continuation when δ is

close to 1. So sufficiency will require a condition such as $v_i(s(t), \delta) > \max_s m_i(s, \delta)$ for all $s(t)$. Call such a strategy *dynamically Individually Rational*. From the previous lemmas, and the fact that $m_i(s, \delta) \rightarrow m_i(s)$, as $\delta \rightarrow 1$, where $m_i(s)$ is the long-run average min-max level for player i the following result is obvious. The min-max limiting result is due to Mertens and Neyman (1983).

Lemma 7 Pick any $v \in \cap_s F(s)$ such that $v_i > \max_s m_i(s)$ for all s . Then there is a time-cycle strategy which is *dynamically Individually Rational* for high δ .

We are now ready to state and prove the main result:

Theorem 3 (Folk Theorem) Suppose that S and A_i are finite sets. Suppose furthermore that T is infinite and that $\cap_s F(s)$ has dimension I (where I is the number of players). Pick any $v \in \cap_s F(s)$ such that $v_i > \max_s m_i(s)$ for all s . Then, for all $\varepsilon > 0$, there is a discount cut-off $\delta(\varepsilon) < 1$ and a time-cycle strategy that for $\delta > \delta(\varepsilon)$ is a SPE with payoffs that are within ε of v .

Proof Without loss of generality, let us set $\max_s m_i(s) = 0$ for all i . From the fact that $\cap_s F(s)$ has dimension I it follows that we can find I payoff vectors in that set – $v^j, j = 1, \dots, I$ – such that for all i (a) $v^j \gg 0$, (b) $v_i^j > v_i^i, j \neq i$, and (c) $v_i > v_i^i$. That we can find these vectors such that (b) is satisfied follows from the dimensionality of the set. That we can additionally get the vectors to satisfy (a) and (c) follows from the fact that it is a convex set and hence an appropriate “averaging” with a vector such as v achieves (a) while an “averaging” with i 's worst payoff achieves (c).

Now consider the following strategy: Norm – Start with a time-cycle strategy that generates payoffs after all histories that are within ε of v . Choose a high enough δ as required. Continue with that strategy if there are no deviations against it. Punishment – If there is, say if player i deviates, then min-max i for T periods and thereafter proceed to the time-cycle strategy that yields payoffs within ε of v^j after all histories. Re-start the punishment whenever there is a deviation.

Choose T in such a fashion that the payoff to the min-max period plus v_i^j is strictly less than v_i . That ensures there is no incentive to deviate against the norm provided the punishment is carried out. That there is incentive for players $j \neq i$ to punish player i follows from the fact that $v_j^i > v_j^j$ the former payoff being what they get from punishing and the latter from not punishing i . That there is incentive for player i not to deviate against his own punishment follows from the fact that re-starting the punishment only lowers his payoffs. The theorem is proved.

A few remarks are in order.

Remark 1 If the game horizon T is finite, there is likely a Folk Theorem along the lines of the result proved for Repeated Games by Benoit and Krishna (1987). To the best of my knowledge it remains, however, an open question.

Remark 2 When the state space A_i is infinite or the state space S is uncountably infinite we again run into technical difficulties. There is an analog to Lemmas 3 and 4 in this instance and under appropriate richer assumptions the results can be generalized – see Dutta (1993). Lemmas 5–7 and the Folk Theorem itself does use the finiteness of S to apply uniform bounds to various approximations and those become problematical when the state space is infinite. It is our belief that nevertheless the Folk Theorem can be proved in this setting. It remains, however, to be done.

Dynamics

Recall that the fourth question is: what can be said about the dynamics of SPE outcome paths? The analogy that might be made is to the various convergence theorems – sometimes also called “turn-pike theorems” – that are known to be true in single-player dynamic programming models. Now even within those models – as has become clear from the literature of the past 20 years in chaos and cycles theory for example – it is not always the case that there are regularities exhibited by the optimal solutions. Matters are worse in dynamic games.

Even within some special models where the single-player optima are well-behaved, the SPE of the corresponding dynamic game need not be. A classic instance is the neo-classical aggregative growth model. In that model, results going back 50 years show that the optimal solutions converge monotonically to a steady-state, the so-called “golden rule” (For references, see Majumdar et al. (2000)). However, examples can be constructed – and may be found in Dutta and Sundaram (1996) and Dockner et al. (1998) – where there are SPE in these models that can have arbitrarily complex state dynamics which for some range of discount factor values descend into chaos. And that may happen with Stationary Markov Perfect Equilibrium. (It would be less of a stretch to believe that SPE in general can have complex dynamics. The Folk Theorem already suggests that it might be so).

There are, however, many questions that remain including the breadth of SPE that have regular dynamics. One may care less for complex dynamic SPE if it can be shown that the “good ones” have regular dynamics. What also remains to be explored is whether adding some noise in the transition equation can remove most complex dynamics SPE.

Global Climate Change – Issues, Models

Issues

The dramatic rise of the world’s population in the last three centuries, coupled with an even more dramatic acceleration of economic development in many parts of the world, has led to a transformation of the natural environment by humans that is unprecedented in scale. In particular, on account of the greenhouse effect, *global warming* has emerged as a central problem, unrivaled in its potential for harm to life as we know it on planet Earth. Seemingly the consequences are everywhere: melting and break-up of the world’s ice-belts whether it be in the Arctic or the Antarctic; heat-waves that set all-time temperature highs whether it be in Western Europe or sub-Saharan Africa; storms increased in frequency and ferocity whether it be Hurricane Katrina or typhoons in

Japan or flooding in Mumbai. In addition to Al Gore's eminently readable book, "An Inconvenient Truth", two authoritative recent treatments are the Stern Review on the Economics of Climate Change, October, 2006 and the IPCC Synthesis Report, November, 2007. Here are three – additional – facts drawn from the IPCC Report:

1. Eleven of the last 12 years (1995–2006) have been amongst the 12 warmest years in the instrumental record of global surface temperatures (since 1850).
2. If we go on with "Business as Usual", by 2100 global sea levels will probably have risen by 9 to 88 cm and average temperatures by between 1.5 and 5.5 °C.

Various factors contribute to global warming, but the major one is an increase in greenhouse gases (GHGs) – primarily, carbon dioxide – so called because they are transparent to incoming shortwave solar radiation but trap outgoing longwave infrared radiation. Increased carbon emissions due to the burning of fossil fuel is commonly cited as the principal immediate cause of global warming. A third relevant fact is:

3. Before the Industrial Revolution, atmospheric CO₂ concentrations were about 270–280 parts per million (ppm). They now stand at almost 380 ppm, and have been rising at about 1.5 ppm annually.

The IPCC Synthesis (2007) says "Warming of the climate system is unequivocal, as is now evident from observations of increases in global average air and ocean temperatures, widespread melting of snow and ice, and rising global average sea level" (IPCC Synthesis Report 2007).

It is clear that addressing the global warming problem will require the coordinated efforts of the world's nations.

In the absence of an international government, that coordination will have to be achieved by way of an international environmental treaty. For a treaty to be implemented, it will have to align the incentives of the signatories by way of rewards for cutting greenhouse emissions and punishments for not doing so. For an adequate analysis of this

problem one needs a dynamic and fully strategic approach. A natural methodology for this then is the theory of Subgame Perfect (Nash) equilibria of dynamic games – which we have discussed at some length in the preceding sections.

Although there is considerable uncertainty about the exact costs of global warming, the two principal sources will be a rise in the sea-level and climate changes. The former may wash away low-lying coastal areas such as Bangladesh and the Netherlands. Climate changes are more difficult to predict; tropical countries will become more arid and less productive agriculturally; there will be an increased likelihood of hurricanes, fires and forest loss; and there will be the unpredictable consequences of damage to the natural habitat of many living organisms. On the other hand, emission abatement imposes its own costs. Higher emissions are typically associated with greater GDP and consumer amenities (via increased energy usage). Reducing emissions will require many or all of the following costly activities: cutbacks in energy production, switches to alternative modes of production, investment in more energy-efficient equipment, investment in R&D to generate alternative sources of energy, etc.

The principal features of the global warming problem are:

- *The Global Common* – although the sources of carbon buildup are localized, it is the total stock of GHGs in the global environment that will determine the amount of warming.
- *Near-irreversibility* – since the stock of greenhouse gases depletes slowly, the effect of current emissions can be felt into the distant future.
- *Asymmetry* – some regions will suffer more than others.
- *Nonlinearity* – the costs can be very nonlinear; a rise in one degree may have little effect but a rise in several degrees may be catastrophic.
- *Strategic Setting* – Although the players (countries) are relatively numerous, there are some very large players, and blocks of like-minded countries, like the US, Western Europe, China, and Japan. That warrants a strategic analysis.

The theoretical framework that accommodates all of these features is an *asymmetric dynamic commons* model with the global stock of greenhouse gases as the (common) state variable. The next sub-section will discuss a few models which have most of the above characteristics.

Models

Before presenting specific models, let us briefly relate the climate change problem to the general dynamic game model that we have seen so far, and provide a historical outline of its study. GHGs form – as we saw above – a global common. The study of global commons is embedded in dynamic commons game (DCG). In such a game the state space S is a single-dimensional variable with a “commons” structure meaning that each player is able to change the (common) state. In particular, the transition function is of the form

$$s(t+1) = q\left(s(t) - \sum_{i=1}^I a_i(t)\right).$$

The first analysis of a DCG may be found in (Levhari and Mirman 1980). That paper considered the particular functional form is which $q\left(s(t) - \sum_{i=1}^I a_i(t)\right) = \left[s(t) - \sum_{i=1}^I a_i(t)\right]^\alpha$ for a fixed fraction α . (And, additionally, Levhari and Mirman assumed the payoffs π_i to be logarithmic). Consequently, the paper was able to derive in closed form a (linear) MPE and was able to analyze its characteristics.

Subsequently several authors – Sundaram (1989), Sobel (1990), Benhabib and Radner (1992), Rustichini (1992), Dutta and Sundaram (1993), Sorger (1998) – studied this model in great generality, without making the specific functional form assumption of Levhari and Mirman, and established several interesting qualitative properties relating to existence of equilibria, welfare consequences and dynamic paths.

More recently in a series of papers by Dutta and Radner on the one hand and Dockner and his co-authors on the other, the DCG model has been directly applied to environmental problems

including the problem of global warming. We shall describe the Dutta and Radner work in detail and also discuss some of the Dockner, Long and Sorger research. In particular, the transition equation is identical in the two models (and described below). What is different is the payoff functions.

We turn now to a simplified climate change model to illustrate the basic strategic ideas. The model is drawn from Dutta and Radner (2008a). In the basic model there is no population growth and no possibility of changing the emissions producing technologies in each country. (Population growth is studied in Dutta and Radner (2006) while certain kinds of technological changes are allowed in Dutta and Radner (2004). These models will be discussed later). However, the countries may differ in their “sizes”, their emissions technologies, and their preferences.

There are I countries. The emission of (a scalar index of) greenhouse gases during period t by country i is denoted by $a_i(t)$. [Time is discrete, with $t = 0, 1, 2, \dots$ ad inf.] Let $A(t)$ denote the global (total) emission during period t ;

$$A(t) = \sum_{i=1}^I a_i(t). \quad (16)$$

The total (global) stock of greenhouse gases (GHGs) at the beginning of period t is denoted by $g(t)$. (Note, for mnemonic purposes we are denoting the state variable – the amount of “gas” – g). The law of motion – or transition function q in the notation above – is

$$g(t+1) = A(t) + \sigma g(t), \quad (17)$$

where σ is a given parameter ($0 < \sigma < 1$). We may interpret $(1 - \sigma)$ as the fraction of the beginning-of-period stock of GHG that is dissipated from the atmosphere during the period. The “surviving” stock, $\sigma g(t)$, is augmented by the quantity of global emissions, $A(t)$, during the same period.

Suppose that the payoff of country i in period t is

$$\pi_i(t) = h_i[a_i(t)] - c_i g(t). \quad (18)$$

The function h_i represents, for example, what country i 's gross national product would be at

different levels of its own emissions, holding the global level of GHG constant. This function reflects the costs and benefits of producing and using energy as well as the costs and benefits of other activities that have an impact on the emissions of GHGs, e. g, the extent of forestation. It therefore seems natural to assume that h_i is a strictly concave C^2 function that reaches a maximum and then decreases thereafter.

The parameter $c_i > 0$ represents the marginal cost to the country of increasing the global stock of GHG. Of course, it is not the stock of GHG itself that is costly, but the associated climatic conditions. As discussed below, in a more general model, the cost would be nonlinear.

Histories, strategies – Markovian strategies – and outcomes are defined in exactly the same way as in the general theory above – and will, hence, not be repeated. Thus associated with each strategy vector σ is a total discounted payoff for each player

$$v_i(\sigma, g_0) \equiv \sum_{t=0}^{\infty} \delta^t \pi_i(t; \sigma, g_0).$$

Similarly, SPE and MPE can be defined in exactly the same way as in the general theory.

The linearity of the model is undoubtedly restrictive in several ways. It implies that the model is unable to analyze catastrophes or certain kinds of feedback effects running back from climate change to economic costs. It has, however, two advantages: first, its conclusions are simple, can be derived in closed-form and can be numerically calibrated; hence may have a chance of informing policy-makers. Second, there is little consensus on what is the correct form of non-linearity in costs. Partly the problem stems from the fact that some costs are not going to be felt for another 50 to 100 years and forecasting the nature of costs on that horizon length is at best a hazardous exercise. Hence, instead of postulating one of many possible non-linear cost functions, all of which may turn out to be incorrect for the long-run, one can opt instead to work with a cost function which may be thought of as a linear approximation to any number of actual non-linear specifications.

Dockner et al. (1998) impose linearity in the emissions payoff function h (whereas in Dutta and Radner it is assumed to be strictly concave) while their cost to g is strictly convex (as opposed to the above specification in which it is linear). The consequent differences in results we will discuss later.

Global Climate Change – Results

In this section we present two sets of results from the Dutta and Radner (2008a) paper. The first set of results characterize two benchmarks – the global Pareto optima, and a simple MPE, called “Business As Usual” and compares them. The second set of results then characterizes the entire SPE correspondence and – relatedly – the best and worst equilibria. Readers are referred to that paper for further results from this model and for a numerical calibration of the model. Furthermore, for the results that are presented, the proofs are merely sketched.

Global Pareto Optima

Let $x = (x_i)$ be a vector of positive numbers, one for each country. A *Global Pareto Optimum (GPO)* corresponding to x is a profile of strategies that maximizes the weighted sum of country payoffs,

$$v = \sum_i x_i v_i, \quad (19)$$

which we shall call *global welfare*. Without loss of generality, we may take the weights, x_i , to sum to 1.

Theorem 4 Let $\hat{V}(g)$ be the maximum attainable global welfare starting with an initial GHG stock equal to g . That function is linear in g ;

$$\begin{aligned} \hat{V}(g) &= \hat{u} - wg, \\ w &= \frac{1}{1 - \delta \sigma} \sum_i x_i c_i, \\ \hat{u} &= \frac{\sum_i x_i h_i(\hat{a}_i) - \delta w \hat{A}}{1 - \delta}. \end{aligned} \quad (20)$$

The optimal strategy is to pick a constant action – emission – every period and after all histories, $\hat{a}_i$ where its level is determined by

$$x_i h'_i(\hat{a}_i) = \delta w. \quad (21)$$

Proof We shall show by dynamic programming arguments that the Pareto-optimal value function is of the form $\hat{V} = \sum_{i=1}^I x_i [\hat{u}_i - w_i g]$. We need to be able to find the constants $\hat{u}_i$ to satisfy:

$$\sum_{i=1}^I x_i [\hat{u}_i - w_i g] = \max_{a_1, \dots, a_I} \sum_{i=1}^I x_i \left[h_i(a_i) - c_i g + \delta \left(\hat{u}_i - w_i \left(\sigma g + \sum_{j=1}^I a_j \right) \right) \right]. \quad (22)$$

Collecting terms that need maximization we can reduce the equation above to

$$\sum_{i=1}^I x_i \hat{u}_i = \max_{a_1, \dots, a_I} \sum_{i=1}^I x_i \left[h_i(a_i) - \delta w_i \sum_{j=1}^I a_j \right] + \delta \sum_{i=1}^I x_i \hat{u}_i. \quad (23)$$

It is clear that the solution to this system is the same for all g ; call this (first-best) solution $\hat{a}_i$. Elementary algebra reveals that

$$\hat{u}_i = \frac{h_i(\hat{a}_i) - \delta w_i \sum_{j=1}^I \hat{a}_j}{1 - \delta} \quad \text{and} \quad w_i = \frac{c_i}{1 - \delta \sigma}.$$

It is also obvious that $x_i h'_i(\hat{a}_i) = \delta w$, where $w = \sum_{i=1}^I x_i w_i$. $\square$

Theorem 4 states that, independently of the level of GHG, g , each country should emit an amount $\hat{a}_i$. The fact that the optimal emission is constant follows from the linearity of the model in g . Notice that on account of the linearity in the gas buildup equation – Eq. 17 – a unit of emission in period t can be analyzed in isolation as a surviving unit of size σ in period $t+1$, σ^2 in period $t+2$, σ^3 in period $t+3$, and so on. On account of the linearity in cost, these surviving units add $(\sum_i x_i c_i) \times \delta \sigma$ in period $t+1$, $(\sum_i x_i c_i) \times (\delta \sigma)^2$ in period $t+2$, and so on, i. e., the marginal lifetime cost is

$$\frac{1}{1 - \delta \sigma} \sum_i x_i c_i,$$

or w , and that marginal cost is independent of g .

A Markov-Perfect Equilibrium: “Business as Usual”

This MPE shares the feature that the equilibrium emission rate of each country is constant in time, and it is the unique MPE with this property. We shall call it the “Business-as-Usual” equilibrium. Note that in this equilibrium each country takes account of the incremental damage to itself caused by an incremental increase in its emission rate, but does not take account of the damage caused to other countries.

Theorem 5 (Business-as-Usual Equilibrium) Let g be the initial stock of GHG. For each country i , let a_i^* be determined by

$$\begin{aligned} h'_i(a_i^*) &= \delta w_i, \\ w_i &= \frac{c_i}{1 - \delta \sigma}, \end{aligned} \quad (24)$$

and let its strategy be to use a constant emission equal to a_i^* in each period; then this strategy profile is a MPE, and country i 's corresponding payoff is

$$\begin{aligned} V_i^*(g) &= u_i^* - w_i g, \\ u_i^* &= \frac{h_i(a_i^*) - \delta w_i A^*}{1 - \delta}. \end{aligned} \quad (25)$$

The intuition for the existence of an MPE with constant emissions is similar to the analogous result for the GPO solution. (And indeed for that reason the proof will be omitted). As long as other countries do not make their emissions contingent on the level of GHGs, country i has a constant marginal lifetime cost to emissions. And that marginal cost is independent of g .

Comparison of the GPO and Business as Usual

The preceding results enable us to compare the emissions in the GPO with those in the Business-as-Usual MPE:

$$\begin{aligned} \text{GPO} : h'_i(\hat{a}_i) &= \frac{\delta \sum_j x_j c_j}{x_i(1 - \delta\sigma)}, \\ \text{BAU} : h'_i(a_i^*) &= \frac{\delta c_i}{1 - \delta\sigma}. \end{aligned} \quad (26)$$

Since

$$x_i c_i < \sum_j x_j c_j,$$

it follows that

$$\frac{\delta c_i}{1 - \delta\sigma} < \frac{\delta \sum_j x_j c_j}{x_i(1 - \delta\sigma)}.$$

Since h_i is concave, it follows that

$$a_i^* > \hat{a}_i. \quad (27)$$

Note that this inequality holds except in the trivial case in which all welfare weights are zero (except one). This result is known as the tragedy of the commons – whenever there is some externality to emissions, countries tend to over-emit in equilibrium. In turn, all this follows from the fact that in the BAU equilibrium each country only considers its own marginal cost and ignores the cost imposed on other countries on account of its emissions; in the GPO solution that additional cost is, of course, accounted for. It follows that the GPO is strictly Pareto superior to the MPE for an open set of welfare weights x_i (and leads to a strictly lower steady-state GHG level for all welfare weights).

One can contrast these results with those in Dockner, Long and Sorger (1998) that studies a model in which the benefits are linear in emission – i. e., h_i is linear – but convex in costs $c_i(\cdot)$. The consequence of linearity in the benefit function h is that the GPO and BAU solutions have a “most rapid approach” (MRAP) property – if $(1 - \sigma)g$, the depreciated stock in the next period, is less than a most preferred g^* , it is optimal to jump the system to g^* . Else it is optimal to wait for depreciation to bring the stock down to g^* . In other words, linearity in benefits implies a “one-shot” move to a desired level of gas g^* , which is thereafter maintained, while linearity in cost (as in the Dutta and Radner model) implies a constant emission rate. What is

unclear in the Dockner, Long and Sorger model is why the multiple players would have the same target steady-state g^* . It would appear natural that, with asymmetric payoffs, each player would have a different steady-state. The existence of a MRAP equilibrium would appear problematical consequently. The authors impose a condition that implies that there is not too much asymmetry.

All SPE

We now turn to the second set of results – a full characterization of SPE in Dutta and Radner (2008a). We will show that the SPE payoff correspondence has a surprising simplicity; the *set* of equilibrium payoffs at a level g is a simple linear translate of the set of equilibrium payoffs from some benchmark level, say, $g = 0$. Consequently, it will be seen that the *set* of emission levels that can arise in equilibrium from level g is identical to those that can arise from equilibrium play at a GHG level of 0. Note that the fact that the set of equilibrium possibilities is invariant to the level of g is perfectly consistent with the possibility that, in a particular equilibrium, emission levels vary with g . However, the invariance property will make for a particularly simple characterization of the best and worst equilibria.

Let $\Xi(g)$ denote the *set* of equilibrium payoff vectors with initial state g , i. e., each element of $\Xi(g)$ is the payoff to *some* SPE starting from g .

Theorem 6 The equilibrium payoff correspondence Ξ is linear; there is a compact set $U \subset \mathcal{R}^I$ such that for every initial state g

$$\Xi(g) = U - \{w_1 g, w_2 g, \dots, w_I g\}$$

where $w_i = c_i/(1 - \sigma\delta)$, $i = 1, \dots, I$. In particular, consider any SPE, any period t and any history of play up until t . Then the payoff vector for the continuation strategies must necessarily be of the form

$$v - (w_1 g_t, w_2 g_t, \dots, w_I g_t).$$

The theorem is proved by way of a bootstrap argument. We presume that a (candidate) payoff

set has this invariance and show that the linear structure of the model confirms the conjecture. Consequently, we generate another candidate pay-off set – which is also state-invariant. Then we look for a fixed point of that operator. In other words, we employ the APS operator to generate the SPE correspondence. Since that has already been discussed in the previous section, it is skipped here.

We will now use the above result to characterize the best – and the worst – equilibria in the global climate change game. Consider the *second-best problem* (from initial state g and for a given vector of welfare weights $\mathbf{x} = (x_i; i = 1, \dots, I)$), i. e., the problem of maximizing a weighted sum of *equilibrium payoffs*:

$$\max \sum_{i=1}^I x_i V_i(g), V(g) \in \Xi(g).$$

Note that we consider all possible equilibria, i. e., we consider equilibria that choose to condition on current and past GHG levels as well as equilibria that do not. The result states that the best equilibrium *need not* condition on GHG levels:

Theorem 7 There exists a constant emission level $a \equiv a_1, a_2, \dots, a_I$ – such that no matter what the initial level of GHG, the second-best policy is to emit at the constant rate a . In the event of a deviation from this constant emissions policy by country i , play proceeds to i 's worst equilibrium. Furthermore, the second-best emission rate is always strictly lower than the BAU rate, i. e., $\bar{a} < a^*$. Above a critical discount factor (less than 1), the second-best rate coincides with the GPO emission rate $\hat{a}$.

The theorem is attractive for three reasons: first, it says that the best possible equilibrium behavior is no more complicated than BAU behavior; so there is no argument for delaying a treaty (to cut emissions) merely because the status quo is simple. Second, the cut required to implement the second-best policy is an across the board cut – independently of anything else, country i should cut its emissions by the amount $a_i^* - \bar{a}_i$.

Third, the second-best is exactly realized at high discount factors, rather than asymptotically approached as the discount factor tends to 1.

Sanctions will be required if countries break with the second-best policy and without loss of generality we can restrict attention to the worst such sanction. We turn now to a characterization of this worst equilibrium (for, say, country i). One definition will be useful for this purpose:

Definition 1 An i -less second-best equilibrium is the solution to a second-best problem in which the welfare weight of i is set equal to zero, i. e., $x_i = 0$.

By the previous theorem, every such problem has a solution in which on the equilibrium path, emissions are a constant. Denote that emission level $a(x_{-i})$:

Theorem 8 There exists a “high” emission level $\bar{a}(i)$ (with $\sum_{j \neq i} \bar{a}_j(i) > \sum_{j \neq i} a_j^*$) and an i -less second-best equilibrium $a(x_{-i})$ such that country i 's worst equilibrium is:

1. Each country emits at rate $\bar{a}_j(i)$ for one period (no matter what g is), $j = 1, \dots, I$.
2. From the second period onwards, each country emits at the constant rate $a_j(x_{-i})$, $j = 1, \dots, I$.

And if any country k deviates at either stages 1 or 2, play switches to k 's worst equilibrium from the very next period after the deviation.

Put another way, for every country i , a sanction is made up of two emission rates, $\bar{a}(i)$ and $a(x_{-i})$. The former imposes immediate costs on country i . The way it does so is by increasing the emission levels of countries $j \neq i$. The effect of this is a temporary increase in incremental GHG but due to the irreversibility of gas accumulation, a permanent increase in country i 's costs, enough of an increase to wipe out any immediate gains that the country might have obtained from the deviation. Of course this additional emission also increases country j 's costs. For the punishing countries, however, this increase is offset by the subsequent permanent change, the switch to the emission vector $\mathbf{a}(x_{-i})$, which permanently increases their quota at the expense of country i 's.

Generalizations

The models discussed thus far are base-line models and do not deal with two important issues relating to climate change – technological change and capital accumulation. Technological change is important because that opens access to technologies that do not currently exist, technologies that may have considerably lower “emissions to energy” ratios, i. e., cleaner technologies. Capital accumulation is important because an important question is whether or not curbing GHGs is inimical to growth. The position articulated by both developing countries like Indian and China as well as by developed economies like the United States is that it is: placing curbs on emissions would restrict economic activities and hence restrain the competitiveness of the economy.

In Dutta and Radner (2004; 2008b) the following modification was made to the model studied in the previous section. It was presumed that the actual emission level associated with energy usage e_i is $f_i e_i$ where f_i is an index of (un) cleanliness – or *emission factor* – higher values implying larger emissions for the same level of energy usage. It was presumed that the emission factor could be changed at cost but driven no lower than some minimum μ_i . In other words,

$$0 \leq e_i(t), \quad (28)$$

$$\mu_i \leq f_i(t+1) \leq f_i(t). \quad (29)$$

Capital accumulation and population growth is also allowed in the model but taken to be exogenous. The dynamics of those two variables are governed by:

$$g(t) = \sigma g(t-1) + \sum_{i=1}^I f_i(t) e_i(t), \quad (30)$$

$$K_i(t+1) = H[K_i(t)], \quad K_i(t) \nearrow \text{and unbounded in } t, \quad (31)$$

$$P_i(t+1) = \psi_i P_i(t) + (1 - \psi_i) \Psi, \quad P_i(t) \leq \Psi. \quad (32)$$

The output (gross-domestic product) of country i in period t is

$$h_i[K_i(t), P_i(t), e_i(t)],$$

where the function h_i has all of the standard properties mentioned above. The damage due to the stock of GHG, $g(t)$, is assumed to be (in units of GDP):

$$c_i P_i(t) g(t).$$

The cost of reducing the emission factor from $f_i(t)$ to $f_i(t+1)$ is assumed to be:

$$\varphi_i[f_i(t) - f_i(t+1)].$$

Immediately it is clear that the state variable now encompasses not just the common stock g but, additionally, the emission factor profile as well as the sizes of population and capital stock. In other words, $s = (g, f, K, P)$. Whilst this significant increase in dimensionality might suggest that it would be difficult to obtain clean characterizations, the papers show that there is some separability. The MPE “Business as Usual” has a separable structure – energy usage $e_i(t)$ and emission factor choice $f_i(t+1)$ – depend solely on country i ’s capital stock and population alone. It varies by period – unlike in the base-line model discussed above – as the exogenous variables vary. Furthermore, the emission factor $f_i(t+1)$ stays unchanged till the population and capital stock cross a threshold level beyond which the cleanest technology gets picked. (This bang-bang character follows from the linearity of the model).

The Global Pareto Optimal solution has similar features – the energy usage in country i is directly driven by the capital stock and population of that country. Furthermore the emission factor choice follows the same bang-bang character as for the MPE. However, there is a tragedy of the common in that in the MPE (versus the Pareto optimum) the energy usage is higher – at every state – and the switch to the cleanest technology happens later.

Future Directions

Within the general theory of dynamic games there are several open questions and possible

directions for future research to take. On the existence question, there needs to be a better resolution of the case where the state space S is uncountably infinite. This is not just a technical curiosity. In applications, typically, in order to apply calculus techniques, we take the state variable to be a subset of some real space. The problem is difficult but one hopes that ancillary assumptions – such as concavity and monotonicity – will be helpful. These assumptions come “cheaply” because they are routinely invoked in economic applications.

The characterization result via APS techniques has a similar technical difficulty blocking its path, as the existence question. The folk theorem needs to be generalized as well to the S infinite case. Here it is our belief though that the difficulty is not conceptual but rather one where the appropriate result needs to be systematically worked out. As indicated above, the study of the dynamics of SPE paths is in its infancy and much remains to be done here.

Turning to the global climate change application, this is clearly a question of utmost social importance. The subject here is very much in the public consciousness yet academic study especially within economics is only a few years old. Many questions remain: generalizing the models to account for technological change and endogenous capital accumulation, examination of a carbon tax, of cap and trade systems for emission permits, of an international bank that can selectively foster technological change, ... There are – as should be immediately clear – enough interesting important questions to exhaust many dissertations and research projects!

Bibliography

- Abreu D (1988) On the theory of infinitely repeated games with discounting. *Econometrica* 56:383–396
- Abreu D, Pearce D, Stachetti E (1990) Towards a general theory of discounted repeated games with discounting. *Econometrica* 58:1041–1065
- Benhabib J, Radner R (1992) The joint exploitation of a productive asset: A game-theoretic approach. *Econ Theory* 2:155–190
- Benoit J-P, Krishna V (1987) Finitely repeated games. *Econometrica* 53:905–922
- Dockner E, Nishimura K (1999) Transboundary boundary problems in a dynamic game model. *Jpn Econ Rev* 50:443–456
- Dockner E, Long N, Sorger G (1996) Analysis of Nash equilibria in a class of capital accumulation games. *J Econ Dyn Control* 20:1209–1235
- Duffie D, Geanakoplos J, Mas-Colell A, McLennan A (1994) Stationary Markov equilibria. *Econometrica* 62(4):745–781
- Dutta P (1991) What do discounted optima converge to? A theory of discount rate asymptotics in economic models. *J Econ Theory* 55:64–94
- Dutta P (1995) A folk theorem for stochastic games. *JET* 66:1–32
- Dutta P, Radner R (2004) Self-enforcing climate change treaties. *Proc Nat Acad Sci USA* 101(14):5174–5179
- Dutta P, Radner R (2006) Population growth and technological change in a global warming model. *Econ Theory* 29:251–270
- Dutta P, Radner R (2008a) A strategic model of global warming model: Theory and some numbers. *J Econ Behav Organ* (forthcoming)
- Dutta P, Radner R (2008b) Choosing cleaner technologies: Global warming and technological change (in preparation)
- Dutta P, Sundaram R (1993) How different can strategic models be? *J Econ Theory* 60:42–61
- Fudenberg D, Maskin E (1986) The Folk theorem in repeated games with discounting or incomplete information. *Econometrica* 54:533–554
- Harris C, Reny P, Robson A (1995) The existence of subgame perfect equilibrium in continuous games with almost perfect information: a case for extensive-form correlation. *Econometrica* 63:507–544
- Inter-Governmental Panel on Climate Change (2007) Climate change, the synthesis report. IPCC, Geneva
- Levhari D, Mirman L (1980) The great fish war: an example using a dynamic cournot-Nash solution. *Bell J Econ* 11:322–334
- Long N, Sorger G (2006) Insecure property rights and growth: the role of appropriation costs, wealth effects and heterogeneity. *Economic Theory* 28:513–529
- Mertens, Neyman (1983)
- Mertens J-F, Parthasarathy T (1987) Equilibria for Discounted Stochastic Games. Research Paper 8750, CORE. University Catholique de Louvain
- Nowak A (1985) Existence of equilibrium stationary strategies in discounted noncooperative stochastic games with uncountable state space. *J Optim Theory Appl* 45:591–603
- Parthasarathy T (1973) Discounted, positive and non-cooperative stochastic games. *Int J Game Theory*:2–1
- Rieder U (1979) Equilibrium plans for non-zero sum Markov games. In: Moeschlin O, Pallasche D (eds) *Game theory and related topics*. North-Holland, Amsterdam

- Rustichini A (1992) Second-best equilibria for games of joint exploitation of a productive asset. *Economic Theory* 2:191–196
- Shapley L (1953) Stochastic games. In: *Proceedings of national academy of sciences*, Jan 1953
- Sobel M (1990) Myopic solutions of affine dynamic models. *Oper Res* 38:847–853
- Sorger G (1998) Markov-perfect Nash equilibria in a class of resource games. *Econ Theory* 11:79–100
- Stern N (2006) Review on the economics of climate change. HM Treasury, London. www.sternreview.org.uk
- Stern Review on the Economics of Climate Change Oct (2006)
- Sundaram R (1989) Perfect equilibrium in a class of symmetric dynamic games. *J Econ Theory* 47:153–177

Static Games

Oscar Volij

Ben-Gurion University, Beer-Sheva, Israel

Article Outline

Glossary

Definition of the Subject

Introduction

Nash Equilibrium

Analysis of Some Finite Games

Existence

Mixed Strategies

The War of Attrition (cont.)

Equilibrium in Beliefs

Correlated Equilibrium

Rationality, Correlated Equilibrium and
Equilibrium in Beliefs

Rationality and Correlated Equilibrium

Bayesian Games

The Asymmetric Information Version of the War
of Attrition

Evolutionary Stable Strategies

Future Directions

Bibliography

Glossary

Player A participant in a game

Action set The set of actions that a player may
choose

Action profile A list of actions, one for each
player

Payoff The utility a player obtains from a given
action profile

Definition of the Subject

Game theory concerns the interaction of decision makers. This interaction is modeled by means of games. There are various approaches to

constructing games. One approach is to focus on the possible outcomes of the decision-makers' interaction by abstracting from the actions or decisions that may lead to these outcomes. The main tool used to implement this approach is the *cooperative game*. Another approach is to focus on the actions that the decision-makers can take, the main tool being the *non-cooperative game*. Within this approach, strategic interactions are modeled in two ways. One is by means of *dynamic*, or extensive form games, and the other is by means of *static*, or strategic games. Dynamic games stress the sequentiality of the various decisions that agents can make. An essential component of a dynamic game is the description of who moves first, who moves second, etc. Static games, on the other hand, abstract from the sequentiality of the possible moves, and model interactions as simultaneous decisions, where the decisions may well be complicated plans of actions that dictate different moves for different situations that may arise. All extensive form games can be modeled as static games, and all strategic form games can be modeled as extensive form games. But some situations may be more conveniently modeled as one or the other kind of game.

This chapter reviews the main ideas and results related to static games, as well as some interesting relationships that connect equilibrium concepts with the idea of rationality. The objective is to introduce the reader to the area of static games and to stimulate his interest for further knowledge of game theory in general. For a comprehensive exposition of some results not covered in this chapter, the reader is referred to the many excellent textbooks available on game theory. Binmore (2007), Fudenberg and Tirole (1991), Osborne (2004), Osborne and Rubinstein (1994) constitute only a partial list.

Although the definition of a static game is a very simple one, static games are a very flexible model which allows us to analyze many different situations. In particular, one can use them to analyze strategic interactions that involve either common interests or diametrically opposed interests. Similarly, one can also use static games to model

situations where players have either symmetric or asymmetric information. The range of applications of static games is very wide and covers many disciplines, such as economics, political science, biology, philosophy, and computer science among others.

Introduction

In this section we introduce some examples that will be used later to motivate different concepts. We also introduce the definition of a static game.

The prisoner's dilemma involves a donor who is interested in donating some amount of money to two universities. The donor decides that the amount each university will receive depends on the content of the messages the presidents of the respective universities will send to him. Each university will send simultaneously one of two messages. One possible message is "Give him 2" and the other is "Give me 1" The donor will do exactly as told. For instance, if University I sends the message "Give me 1" and University II sends "Give him 2", the donor will donate \$3 to University I and \$0 to University II. This game can be described by means of the following matrix, where the entries represent the payoffs for University I and University II, respectively, that result from the corresponding action choices.

		University II	
		Give him 2	Give me 1
University I	Give him 2	2, 2	0, 3
	Give me 1	3, 0	1, 1

The battle of the sexes consists of two friends, She and He, who want to go out together, but have no means of communication. They have to decide, each one separately but both simultaneously, whether to go to a boxing match or to a ballet show. For both of them, the worst possible outcome would be to choose different events and not meet. But if they meet, he would rather meet her at the boxing match, while she would rather meet him at the ballet. The battle of the sexes can be described by the following matrix.

		She	
		Box	Ballet
He	Box	2, 1	0, 0
	Ballet	0, 0	1, 2

Again, the entries of this matrix represent the payoffs that he and she get, as a result of their corresponding choices.

Chicken models two drivers who approach each other on a narrow street. If none of them slows down they will have an accident and their corresponding payoffs will be 0. But if at least one of them slows down, the accident is prevented. The problem is that both of them would like the other to slow down. If only one driver slows down, this driver gets a payoff of 2 and the other driver gets a payoff of 7. If both drivers slow down, then both drivers get a payoff of 6. This situation can be described by the following matrix.

		Driver 2	
		Slow Down	Speed up
Driver 1	Slow Down	6, 6	2, 7
	Speed up	7, 2	0, 0

Matching Pennies involves two friends, each of whom places a coin on a table. If both coins are placed heads up or tails up, then friend 1 gets one dollar from friend 2. If one coin is placed heads up and the other tails up, then friend 1 pays one dollar to friend 2. Matching pennies can be described by the following matrix, where the entries are the amounts of money that the friends get from each other.

		Friend 2	
		Heads	Tails
Friend 1	Heads	1, -1	-1, 1
	Tails	-1, 1	1, -1

The above examples of strategic interactions can be modeled as static games. A static game is a formalization of a strategic situation according to which players choose their actions separately and simultaneously, and as a result obtain certain payoffs. The interaction that a static game models need not require that players take their actions

simultaneously. But the interaction is modeled by defining actions in such a way that lets us think of the players as acting simultaneously.

All of the above examples involve a set of players, and for each player there is a set of available actions and a function that associates a payoff level to each of the profiles of actions that may result from the players' choices. These are the three essential components of a static game, as formalized in the following definition.

Definition 1 A static game is a triple $\langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ where N is a finite set of players, and for each player $i \in N$, A_i is i 's set actions, and $u_i : \times_{k \in N} A_k \rightarrow \mathbb{R}$ is player i 's utility function.

In the prisoner's dilemma the set of players is $N = \{\text{University I, University II}\}$; the sets of actions are $A_I = A_{II} = \{\text{Give me 1, Give him 2}\}$; the utility function of University I is $u_I(\text{Give me 1, Give me 1}) = 1$, $u_I(\text{Give me 1, Give him 2}) = 3$, $u_I(\text{Give him 2, Give me 1}) = 0$, $u_I(\text{Give him 2, Give him 2}) = 2$; and the utility function of University II is $u_{II}(\text{Give me 1, Give me 1}) = 1$, $u_{II}(\text{Give me 1, Give him 2}) = 0$, $u_{II}(\text{Give him 2, Give me 1}) = 3$, $u_{II}(\text{Give him 2, Give him 2}) = 1$.

In this chapter we sometimes refer to static games simply as games. For any game $\langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$, the set of action profiles $\times_{k \in N} A_k$ is denoted by A , and a typical action profile is denoted by $a = (a_i)_{i \in N} \in A$. If A is a finite set, then we say that the game is finite. Player i 's utility function represents his preferences over the set of action profiles. For instance, for any two action profiles a and a' in A , $u_i(a) \geq u_i(a')$ means that player i prefers action profile a to action profile a' . Clearly, although player i has preferences over action profiles, he can only affect his own component, a_i , of the profile.

Nash Equilibrium

One objective of game theory is to select, for each game, a set of action profiles that are interesting in some way. These action profiles may be interpreted as predictions of the theory, or prescriptions for the

players to follow, or simply as equilibrium outcomes in the sense that if they occur, the players do not wish that they had acted differently. These action profiles are formally given by solution concepts, which are functions that associate each strategic game with the selected set of action profiles. The central solution concept in game theory is known as *Nash equilibrium*. The hypothesis behind this solution concept is that each player chooses his actions so as to maximize his utility, given the profile of actions chosen by the other players. To give a formal definition of the Nash equilibrium concept, we first introduce some useful notation. For each player $i \in N$, let $A_{-i} = \times_{k \in N \setminus \{i\}} A_k$ be the set of the other players' profiles of actions. Then we can write $A = A_i \times A_{-i}$, and each action profile can be written as $a = (a_i, a_{-i}) \in A_i \times A_{-i}$, thereby distinguishing player i 's action from the other players' profile of actions.

Definition 2 The action profile $a^* = (a_i^*)_{i \in N} \in A$ in a game $\langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$, is a *Nash equilibrium* if for each player, $i \in N$, and every action $a_i \in A_i$ of player i , a^* is at least as good for player i as the action profile (a_i, a_{-i}^*) . That is, if

$$u_i(a^*) \geq u_i(a_i, a_{-i}^*) \text{ for all } a_i \in A_i \text{ and for all } i \in N.$$

It is a *strict Nash equilibrium* if the above inequality is strict for all alternative actions $a_i \in A_i \setminus \{a_i^*\}$.

Analysis of Some Finite Games

Prisoner's Dilemma Recall that the prisoner's dilemma can be described by the following matrix.

		University II	
		Give him 2	Give me 1
University I	Give him 2	2, 2	0, 3
	Give me 1	3, 0	1, 1

The action profile $(\text{Give me 1, Give me 1})$ is a Nash equilibrium. Indeed,

$$\begin{aligned}
u_I(\text{Give me 1, Give me 1}) &= 1 \\
&\geq u_I(\text{Give me 2, Give me 1}) = 0
\end{aligned}$$

and

$$\begin{aligned}
u_{II}(\text{Give me 1, Give me 1}) &= 1 \\
&\geq u_{II}(\text{Give me 1, Give me 2}) = 0.
\end{aligned}$$

On the other hand, the action profile (*Give him 2, Give him 2*) is not a Nash equilibrium, since University I prefers action “*Give me 1*” if University II chooses action “*Give him 2*”:

$$\begin{aligned}
2 &= u_I(\text{Give him 2, Give him 2}) \\
&< u_I(\text{Give him 1, Give him 2}) = 3
\end{aligned}$$

Battle of the Sexes Recall that the battle of the sexes can be described by the following matrix.

		She	
		Box	Ballet
He	Box	2, 1	0, 0
	Ballet	0, 0	1, 2

One can check that (*Box, Box*) is a Nash equilibrium and (*Ballet, Ballet*) is a Nash equilibrium as well. It can also be checked that these are the only two action profiles that constitute a Nash equilibrium.

Matching Pennies The reader can check that Matching Pennies has no Nash equilibrium.

Before we analyze the next example, we introduce a technical tool that allows us to reformulate the definition of Nash equilibrium more conveniently. More importantly, this alternative definition is the key to the standard proof of the existence of Nash equilibrium.

Definition 3 Let $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ be a strategic game and let $i \in N$ be a player. Consider a list of actions $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n) \in \times_{k \in N \setminus \{i\}} A_k$ of all the players other than i . The set of player i 's best responses to a_{-i} is

$$\begin{aligned}
\mathcal{B}_i(a_{-i}) &= \{a_i \in A_i : u_i(a_i, a_{-i}) \geq u_i(b_i, a_{-i}) \\
&\quad \text{for all } b_i \in A_i\}.
\end{aligned}$$

The correspondence $\mathcal{B}_i : x_j \neq i A_j \rightarrow A_i$ that assigns to each $(n - 1)$ -tuple of actions in A_{-i} the set of best responses to it is called the *best response correspondence* of player i .

The definition of a Nash equilibrium may be stated in terms of the players' best response correspondences, as stated in the following proposition.

Proposition 1 *The action profile $a^* \in A$ is a Nash equilibrium if and only if every player's action is a best response to the other players' actions. That is, if*

$$a_i^* \in \mathcal{B}_i(a_{-i}^*) \quad \text{for all } i \in N.$$

Until now, all the examples involved games where the action sets contained two actions. The next example is a game where the players' action sets are infinite. We will use the player's best response correspondences to find all its Nash equilibria.

The War of Attrition Two animals, 1 and 2, are fighting over a prey. Each animal chooses a time at which it intends to give up. Once one animal has given up, the other obtains the prey; if both animals give up at the same time then they split the prey equally. For each $i = 1, 2$, animal i 's willingness to fight for the prey is given by $v_i > 0$. The value v_i is the maximum amount of time that animal i is willing to spend to obtain the prey. Since fighting is costly, each animal prefers as short a fight as possible. If animal i obtains the prey after a fight of length t , his utility will be $v_i - t$. We can model the situation as the game, $(G = \langle \{1, 2\}, (A_1, A_2), (u_1, u_2) \rangle)$ where

- $A_1 = [0, \infty] = A_2$ (an element $t \in A_i$ represents a time at which player i plans to give up)
- $u_1(t_1, t_2) = \begin{cases} -t_1 & \text{if } t_1 < t_2 \\ \frac{1}{2}v_1 - t_2 & \text{if } t_1 = t_2 \\ v_1 - t_2 & \text{if } t_1 > t_2 \end{cases}$
- $u_2(t_1, t_2) = \begin{cases} -t_2 & \text{if } t_2 < t_1 \\ \frac{1}{2}v_2 - t_1 & \text{if } t_1 = t_2 \\ v_2 - t_1 & \text{if } t_2 > t_1 \end{cases}$

We are interested in the best response correspondences. First, we calculate player 1's best response correspondence, $\mathcal{B}_1(t_2)$. There are three cases to consider.


- Case 1** $t_2 < v_1$ In this case, $v_1 - t_2 > \frac{1}{2}v_1 - t_2$ and $v_1 - t_2 > -t_1$. Consequently, given that player 2's action is t_2 , player 1's utility function has a maximum value of $v_1 - t_2$, which is attained at any $t_1 > t_2$. Therefore, $\mathcal{B}_1(t_2) = (t_2, \infty)$.
- Case 2** $t_2 = v_1$ In this case, $0 = v_1 - t_2 > \frac{1}{2}v_1 - t_2$. Therefore, player's 1 utility function $u_1(\cdot, t_2)$ has a maximum value of 0, which is attained at $t_1 = 0$ and at $t_1 > t_2$. Therefore, $\mathcal{B}_1(t_2) = \{0\} \cup (t_2, \infty)$.
- Case 3** $t_2 > v_1$ In this case $\frac{1}{2}v_1 - t_2 < v_1 - t_2 < 0$. As a result, player 1's utility function $u_1(\cdot, t_2)$ has a maximum value of 0, which is attained at $t_1 = 0$. Therefore, $\mathcal{B}_1(t_2) = \{0\}$

Summarizing, player 1's best response correspondence is:

$$\mathcal{B}_1(t_2) = \begin{cases} (t_2, \infty) & \text{if } t_2 < v_1 \\ \{0\} \cup (t_2, \infty) & \text{if } t_2 = v_1 \\ \{0\} & \text{if } t_2 > v_1 \end{cases}$$

which is depicted in Fig. 1.

Similarly, player 2's best response correspondence is:




Static Games, Fig. 1 Player 1's best response correspondence

$$\mathcal{B}_2(t_1) = \begin{cases} (t_1, \infty) & \text{if } t_1 < v_2 \\ \{0\} \cup (t_1, \infty) & \text{if } t_1 = v_2 \\ \{0\} & \text{if } t_1 > v_2 \end{cases}$$

Combining the two best response correspondences we get that (t_1^*, t_2^*) is a Nash equilibrium if and only if either $t_1^* = 0$ and $t_2^* \geq v_1$ or $t_2^* = 0$ and $t_1^* \geq v_2$. Figure 2 depicts the set of all the Nash equilibria as the intersection of the two best response correspondences. Two things are worth noting. First, it is not necessarily the case that the player who values the prey most wins the war. That is, there are Nash equilibria of the war of attrition where the player with the highest willingness to fight for the prey gives in first, and as a result the object goes to the other player. Second, in none of the Nash equilibria is there a physical fight. All Nash equilibria involve one player giving in immediately to the other. This second feature seems rather unrealistic, since fights in war of attrition-like situations are commonly observed. If one wants to obtain a fight of positive length in the war of attrition one needs to either drop the Nash equilibrium concept and adopt an alternative one, or model the war of attrition differently. We will adopt this second course of action later.

Existence

As the matching pennies example shows, not all games have a Nash equilibrium. The following theorem, which dates back to Nash (1950) and Glicksberg (1952), states sufficient conditions on a game for it to have a Nash equilibrium. An earlier



Static Games, Fig. 2 The equilibria

version of this theorem for the smaller but prominent class of zero-sum games can be found in von Neumann (1928) (translated in von Neumann (1959)). The standard proofs use Kakutani's fixed point theorem. We present here an alternative proof, due to Geanakoplos (2003), which uses Brouwer's fixed point theorem instead.

Theorem 1 *The game $\langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ has a Nash equilibrium if for all $i \in N$*

- *the set A_i of actions of player i is a nonempty compact convex subset of an Euclidean space,*
- *the utility function u_i is continuous,*
- *the utility function u_i is concave in A_i .*

Proof (Geanakoplos) Define the correspondence $\varphi_i: A \rightarrow A_i$ by

$$\varphi_i(\bar{a}) = \arg \max_{a_i \in A_i} \{U_i(a_i, \bar{a}_{-i}) - \|a_i - \bar{a}_i\|^2\},$$

where, $\|\cdot\|$ denotes a norm in the relevant Euclidean space. Note first that φ_i is a nonempty valued correspondence because the maximand is a continuous function and A_i is compact. Second, note that the function $\|a_i - \bar{a}_i\|$ is convex:

$$\begin{aligned} & \|(\lambda a_i + (1 - \lambda)b_i) - \bar{a}_i\| \\ &= \|(\lambda a_i - \lambda \bar{a}_i) + ((1 - \lambda)b_i - (1 - \lambda)\bar{a}_i)\| \\ &\leq \|(\lambda a_i - \lambda \bar{a}_i)\| + \|((1 - \lambda)b_i - (1 - \lambda)\bar{a}_i)\| \\ &\leq \lambda \|a_i - \bar{a}_i\| + (1 - \lambda) \|b_i - \bar{a}_i\| \end{aligned}$$

Since the quadratic function is strictly convex, then the maximand is a strictly concave function. Therefore, the correspondence φ_i is in fact a function. Furthermore, since the maximand is continuous in the parameter $\bar{a}$, φ_i is also continuous. To see this, let $\bar{a}_n \rightarrow \bar{a}$ be a convergent sequence of action profiles and let $a_{in} = \varphi_i(\bar{a}_n)$. This means that $U(a_{in}, (\bar{a}_n)_{-i}) \geq U(b_i, (\bar{a}_n)_{-i})$, for all $b_i \in A_i$. Since A_i is a compact set, a_{in} has a convergent subsequence. Denoting by a_i the limit of this subsequence and applying limits to the above inequality, we obtain that

$$U(a_i, \bar{a}_{-i}) \geq U(b_i, \bar{a}_{-i}) \quad \text{for all } b_i \in A_i,$$

namely $a_i = \varphi_i(\bar{a})$. Since this is true for every convergent subsequence of a_{in} , we have that

$\varphi_i(\bar{a}_n) = a_{in} \rightarrow a_i = \varphi_i(\bar{a})$, which means that φ is continuous.

Now define $\varphi: A \rightarrow A$ by $\varphi = (\varphi_1, \dots, \varphi_N)$. Clearly, φ is a continuous function mapping a compact set to itself. Therefore, by Brouwer's fixed point theorem, it has a fixed point: $\varphi(\bar{a}) = \bar{a}$. We now show that $\bar{a}$ is a Nash equilibrium of the game. Assume not. Then, there is some $i \in N$ with $a_i \in A_i$ such that $U_i(a_i, \bar{a}_{-i}) - U_i(\bar{a}) = E > 0$. Then, by concavity of U_i , for all $0 < \varepsilon < 1$,

$$\begin{aligned} & U_i(\varepsilon a_i + (1 - \varepsilon)\bar{a}_i, \bar{a}_{-i}) - U_i(\bar{a}) \\ &\geq \varepsilon U_i(a_i, \bar{a}_{-i}) + (1 - \varepsilon)U_i(\bar{a}) - U_i(\bar{a}) \\ &\geq \varepsilon E > 0, \end{aligned}$$

while $\|\varepsilon a_i + (1 - \varepsilon)\bar{a}_i - \bar{a}_i\|^2 = \varepsilon^2 \|a_i - \bar{a}_i\|^2 < \varepsilon E$, for small enough ε . Therefore, for such small ε , the action $\varepsilon a_i + (1 - \varepsilon)\bar{a}_i$ satisfies

$$\begin{aligned} & U_i(\varepsilon a_i + (1 - \varepsilon)\bar{a}_i, \bar{a}_{-i}) - \\ & \|\varepsilon a_i + (1 - \varepsilon)\bar{a}_i - \bar{a}_i\|^2 > U_i(\bar{a}) \end{aligned}$$

which contradicts the fact that $\varphi_i(\bar{a}) = \bar{a}_i$. $\square$

Mixed Strategies

So far, we have formally defined a game, and have introduced the solution concept of Nash equilibrium which is arguably the central solution concept of game theory. However, there seem to be two problems with this concept. One is that although Nash equilibria exist in a wide class of games, there are many simple games that do not have a Nash equilibrium. The most troubling example is Matching Pennies. If game theory cannot provide a prediction for this simple game then one must wonder if there is any value to the theory. The second problem is that the concept of Nash equilibrium predicts a very unrealistic outcome in the war of attrition. One would expect that game theory would not only provide nonempty predictions, but also ones that look reasonable and help explain what we see around us.

One way to approach these problems is not to abandon the theory or the concept of Nash

equilibrium altogether, but to modify the way we model the problematic situations. The idea behind mixed strategies is to first modify the game by extending the set of actions available to the players, and then to apply the concept of Nash equilibrium to this extended game. In this way one may obtain additional Nash equilibria, some of which may provide reasonable predictions to the game.

Let $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$, be a finite game. For any A_i , a probability distribution on A_i is a function

$$x_i : A_i \rightarrow \mathbb{R}^+$$

such that

$$\sum_{a_i \in A_i} x_i(a_i) = 1.$$

The set of all probability distributions on A_i is denoted by $\Delta(A_i)$. A *mixed strategy* on A_i is a random choice over elements of A_i , namely an element of $\Delta(A_i)$. If x_i is a mixed strategy on A_i , $x_i(a_i)$ denotes the probability that action $a_i \in A_i$ is selected when x_i is adopted. Since elements of $\Delta(A_i)$ can have an alternative interpretation, such as beliefs about the choice of player i , we denote the set of mixed strategies by X_i to distinguish it from the more abstract set of probability distributions on A_i . Also, we denote the set of mixed strategy problem as $X = \times_{i \in N} X_i$. Denoting for each player $i \in N$, $X_{-i} = \times_{k \in N \setminus \{i\}} X_k$, a typical mixed strategy profile can be written as $(x_k)_{k \in N} = (x_i, x_{-i}) \in X_i \times X_{-i}$. The *mixed extension* of the strategic game G is the strategic game $\langle N, (X_i)_{i \in N}, (U_i)_{i \in N} \rangle$ where the set of actions of player i is the set of mixed strategies, X_i , and the payoff function $U_i : \times_{i \in N} X_i \rightarrow \mathbb{R}$ of player i is defined by

$$U_i((x_k)_{k \in N}) = \sum_{a=(a_k)_{k \in N} \in A} u_i(a) \prod_{k \in N} x_k(a_k).$$

Remark 1 Since each mixed strategy of player i , x_i , can be identified with a vector $x_i = (x_i(a_i))_{a_i \in A_i} \in \mathbb{R}^{|A_i|}$, the function U_i is multinomial in the coordinates of its variables, and, as a result, it is continuous as a function of the players' mixed strategies.

Definition 4 An *equilibrium in mixed strategies* of the game, $\langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ is a Nash equilibrium of the mixed extension of the game. In other words, it is a list of mixed strategies $(x_k^*)_{k \in N} \in X$ such that for all players $i \in N$ and for all his mixed strategies x_i ,

$$U_i((x_k^*)_{k \in N}) \geq U_i((x_i, x_{-i}^*))$$

Alternatively, $(x_k^*)_{k \in N} \in X$ is a mixed strategy equilibrium if

$$x_i^* \in \mathcal{B}_i(x_{-i}^*) \quad \text{for all } i \in N.$$

Note that for every finite game $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$, its mixed extension is a strategic game that satisfies the conditions Theorem 1. As a result, every finite game has a mixed strategy equilibrium.

Example 1 Consider again Matching Pennies. Its mixed extension is the game $\langle N, (X_i)_{i \in N}, (U_i)_{i \in N} \rangle$, where the set of players is $N = \{1, 2\}$, the sets of mixed strategies are $X_1 = \{(p_H, p_T) \geq (0, 0): p_H + p_T = 1\}$, and $X_2 = \{(q_H, q_T) \geq (0, 0): q_H + q_T = 1\}$, and the utility functions are given by $U_1((p_H, p_T), (q_H, q_T)) = p_H q_H + p_T q_T - p_H q_T - p_T q_H$ and $U_2((p_H, p_T), (q_H, q_T)) = p_H q_T + p_T q_H - p_H q_H - p_T q_T$. It can be checked that the only Nash equilibrium of this mixed extension is $((1/2, 1/2), (1/2, 1/2))$. Indeed, since $U_1((p_H, p_T), (1/2, 1/2))$ is identically 0, it attains its maximum at, among other strategies, $(1/2, 1/2)$. The same is true for $U_2((1/2, 1/2), (q_H, q_T))$. To see that there is no other equilibrium, note that for (q_H, q_T) with $q_H > q_T$ player 1's best response is $(0, 1)$. But player 2's best response to $(0, 1)$ is $(1, 0)$. Since $0 \leq 1$, (q_H, q_T) with $q_H > q_T$ cannot be part of an equilibrium. Similarly, for any (q_H, q_T) with $q_H < q_T$ player 1's best response is $(0, 1)$. But player 2's best response to $(0, 1)$ is $(1, 0)$. Since $1 \geq 0$, (q_H, q_T) with $q_H < q_T$ cannot be part of an equilibrium.

We next present a characterization of the mixed strategy equilibria of a game that will sometimes

allow us to compute them in an easy way. Further, this characterization serves as the basis of an interesting interpretation of the mixed strategy equilibrium concept that we will discuss later. For this purpose, we identify the action $a_i \in A_i$ of player i with the mixed strategy of player i that assigns probability 1 to action a_i , and 0 to all other actions. Therefore, given a player i , one of his actions $a_i \in A_i$, and a profile $x = (x_k)_{k \in N}$ of the players' mixed strategies, (a_i, x_{-i}) denotes the mixed strategy profile obtained from x by replacing i 's mixed strategy x_i by the mixed strategy of player i that assigns probability 1 to action a_i . With this notation we can state the following identity:

$$U_i((x_k)_{k \in N}) = \sum_{a_i \in A_i} x_i(a_i) U_i((a_i, x_{-i})). \quad (1)$$

Indeed,

$$\begin{aligned} U_i((x_k)_{k \in N}) &= \sum_{a=(a_k)_{k \in N} \in A} u_i(a) \Pi_{k \in N} x_k(a_k) \\ &= \sum_{a_i \in A_i} \sum_{a_{-i} \in A_{-i}} u_i(a) \Pi_{k \in N} x_k(a_k) \\ &= \sum_{a_i \in A_i} x_i(a_i) \sum_{a_{-i} \in A_{-i}} u_i(a) \Pi_{k \in N \setminus \{i\}} x_k(a_k) \\ &= \sum_{a_i \in A_i} x_i(a_i) U_i((a_i, x_{-i})). \end{aligned}$$

Identity (1) is useful to prove the following characterization of the mixed strategy Nash equilibria.

Lemma 1 *The strategy profile $x^* = (x_k^*)_{k \in N}$ is an equilibrium of the mixed extension of $\langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ if and only if for all players $i \in N$ and for all $a_i \in A_i$,*

$$\text{If } x_i^*(a_i) > 0 \quad \text{then } U_i((a_i, x_{-i}^*)) = U_i(x^*) \quad (2)$$

$$\text{If } x_i^*(a_i) = 0 \quad \text{then } U_i((a_i, x_{-i}^*)) \leq U_i(x^*). \quad (3)$$

Proof Assume that $x^* = (x_k^*)_{k \in N}$ satisfies conditions (2) and (3). Let $i \in N$, and let x_i be a mixed strategy of player i . Then, by (1)

$$\begin{aligned} U_i(x_i, x_{-i}^*) &= \sum_{a_i \in A_i} x_i(a_i) U_i((a_i, x_{-i}^*)) \\ &\leq \sum_{a_i \in A_i} x_i(a_i) U_i(x^*) = U_i(x^*) \end{aligned}$$

and therefore x^* is an equilibrium.

Assume now that $x^* = (x_k^*)_{k \in N}$ is an equilibrium. Let $i \in N$. Then

$$U_i(x^*) \geq U_i((a_i, x_{-i}^*)) \quad \forall a_i \in A_i \quad (4)$$

and, in particular, condition (3) holds for all $a_i \in A_i$ such that $x_i^*(a_i) = 0$. Also, using (1) we can write

$$\sum_{a_i \in A_i} x_i^*(a_i) U_i(x^*) = \sum_{a_i \in A_i} x_i^*(a_i) U_i((a_i, x_{-i}^*)). \quad (5)$$

If there is $a_i \in A_i$ such that $x_i^*(a_i) > 0$ and $U_i(x^*) > U_i((a_i, x_{-i}^*))$ then, using (4),

$$\sum_{a_i \in A_i} x_i^*(a_i) U_i(x^*) > \sum_{a_i \in A_i} x_i^*(a_i) U_i((a_i, x_{-i}^*))$$

in contradiction to (5). $\square$

Corollary 1 *The strategy profile $x^* = (x_k^*)_{k \in N}$ is an equilibrium of the mixed extension of $\langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ and only if for all players $i \in N$ and for all $a_i \in A_i$,*

$$x_i^*(a_i) > 0 \quad \text{implies } a_i \in \mathcal{B}_i(x_{-i}^*).$$

According to the standard interpretation, a player's mixed strategy in a game G is an action, but in a different game, namely in the mixed extension of G . According to this interpretation, a mixed strategy is a deliberate choice of a player to use a random device. A mixed strategy equilibrium then is a profile of independent random devices, each of which is a best response to the others. Corollary 1 provides an alternative interpretation of a mixed strategy equilibrium. According to this interpretation, a player's mixed strategy represents the uncertainty in the minds of the other players concerning the player's action.

In other words, a player's mixed strategy is interpreted not as a deliberate choice of the player but the belief, shared by all the other players, about the player's choice. That is, if $(x_k)_{k \in N}$ is a profile of mixed strategies, then x_i is the conjecture, shared by all the players other than i , about i 's ultimate choice action. Consequently, x_{-i} are the conjectures entertained by player i about his opponents' actions. According to this interpretation, Corollary 1 says that a mixed strategy equilibrium $(x_k^*)_{k \in N}$ is a profile of beliefs about each player's actions (entertained by the other players) according to which each player chooses an action that is a best response to his own beliefs.

The War of Attrition (cont.)

We have seen in section "Analysis of Some Finite Games" that all the Nash equilibria of the war of attrition predict no real fight for the prey. We will now see that there is a mixed strategy equilibrium of the war of attrition that predicts a positive-length fight with probability one.

The players' action sets in the war of attrition are intervals of real numbers. A mixed strategy for player i in that game can be represented by a cumulative distribution function $F_i : [0, \infty] \rightarrow [0, 1]$. For each $t \in (0, \infty]$, $F_i(t)$ is the probability that player i gives up at or before t . We will look for a Nash equilibrium (F_1, F_2) that consists of two strictly increasing, differentiable cumulative distribution functions. The density of F_i is denoted by f_i . We will try to find an equilibrium at which each player is indifferent between all pure actions.

Consider player i . Given that his opponent is using mixed strategy F_j , $j \neq i$, if he chooses to give in at time t , then he will face a lottery according to which,

- With probability $1 - F_j(t)$, player i does not obtain the prey and gets a payoff of $-t$,
- With probability $F_j(t)$, player i obtains the prey at time t_j , where t_j is a random variable whose cumulative distribution function is $F_j(t_j)/F_j(t)$ (the distribution player j 's surrender time, conditional on his having given in before t).

Therefore, the corresponding expected utility of choosing time t is

$$U_i(t, F_j) = (1 - F_j(t))(-t) + F_j(t) \int_0^t (v_i - t_j) d\frac{F_j(t_j)}{F_j(t)} = (1 - F_j(t))(-t) + \int_0^t (v_i - t_j) dF_j(t_j).$$

Since in the equilibrium we are looking for, player i is indifferent among all his actions, the above expression is independent of t . Namely, $U_i(t, F_j) \equiv c$. As a result, the derivative of the above utility with respect to t equals 0. Formally,

$$\begin{aligned} \frac{\partial U_i(t, F_j)}{\partial t} &= t f_j(t) - (1 - F_j(t)) + (v_i - t) f_j(t) \\ &= (1 - F_j(t)) + v_i f_j(t) = 0. \end{aligned}$$

This is a differential equation whose general solution is

$$F_j(t) = 1 - K e^{-\frac{t}{v_i}}$$

If we want it to satisfy $F_j(0) = 0$, we obtain that $K = 1$. As a result, the distribution function is given by

$$F_j(t) = 1 - e^{-\frac{t}{v_i}}$$

Consequently, the equilibrium we are looking for is

$$(F_1(t), F_2(t)) = \left(1 - e^{-\frac{t}{v_2}}, 1 - e^{-\frac{t}{v_1}}\right).$$

According to this equilibrium, for any t , the probability that there is a fight that lasts at least t is $(1 - F_1(t))(1 - F_2(t)) > 0$. Consequently, there is a fight with probability one. The introduction of mixed strategies allowed the concept of Nash equilibrium to be consistent with fights that last a positive length of time. However, the mixed strategy equilibrium has the following unfortunate property. If $v_1 < v_2$, then for all $t > 0$, $F_1(t) < F_2(t)$. In other words, it is more likely that the player with the highest willingness to fight for the prey gives up earlier than any given t , than that the player with the lowest willingness to fight gives in earlier than the same t . Therefore, in equilibrium it is more likely that the player with the lower

willingness to fight wins the war than the other way around. In particular, the probability that player 1 gets the object is given by

$$\int_0^\infty F_2(t) dF_1(t)$$

which can be checked to be equal to $\frac{v_2}{v_1+v_2} > 1/2$.

In order to obtain the more intuitive result that the higher the willingness to fight for the prey, the higher is the probability to obtain it, we will need to model the war of attrition in yet a different way. We'll return to this when we introduce asymmetric information to the games.

Equilibrium in Beliefs

The mixed extension of the game $\langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ is constructed in two steps. First, we enlarge the set of actions available to each player by allowing him to choose any mixed strategy on his original action set. Second, since the action choices are now probability distributions over actions, we extend the players' original preferences to preferences over profiles of mixed strategies. We do so by evaluating each mixed strategy profile according to the expected value of the original utilities with respect to the probability distribution over action profiles induced by the mixed strategy.

The first step seems uncontroversial since it is certainly possible for players to use random devices. But the second step is somewhat problematic because, by evaluating mixed strategies according to the expected utility of the resulting lotteries, one is implicitly imposing on the players a certain kind of risk preferences. One may wonder what the implications would be if instead of extending the preferences by assuming that players are expected utility maximizers, we assume that players have more general preferences over profiles of mixed strategies. In particular, we would like to know if there is a suitable generalization of Corollary 1.

Let $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ be a finite game. We define the mixed extension of G as the

strategic game $\langle N, (X_i)_{i \in N}, (U_i)_{i \in N} \rangle$ where, as in Sect. "Mixed Strategies", X_i is the set of probability distributions over the actions in A_i , for $i \in N$, but unlike there, the utility function $U_i : X \rightarrow \mathbb{R}^N$ is not necessarily a multilinear function of the probabilities, but a general continuous function of the mixed strategies. The only requirement on U_i is that for all profiles of degenerate mixed strategies $(a_k)_{k \in N}$, we have $U_i((a_k)_{k \in N}) = u_i(a_k)_{k \in N}$. As before, a mixed strategy Nash equilibrium of $\langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ is a Nash equilibrium of its mixed extension $\langle N, (X_i)_{i \in N}, (U_i)_{i \in N} \rangle$. In other words, it is a list of mixed strategies $(x_k^*)_{k \in N}$ such that for all players $i \in N$ and for all of his mixed strategies x_i ,

$$U_i((x_k^*)_{k \in N}) \geq U_i((x_i, x_{-i}^*)).$$

Alternatively, $(x_k^*)_{k \in N}$ is a mixed strategy equilibrium if

$$x_i^* \in \mathcal{B}_i(x_{-i}^*) \quad \text{for all } i \in N.$$

Observation 1 It is important to note that two different actions of a player may be best responses to a given mixed strategy profile of the other players, and yet no probability mixture of the two actions will be a best response to the given mixed strategy profile. This will typically be the case when the function U_i is strictly convex in X_i , since strictly convex functions attain their maximum at boundary points.

Theorem 1 shows that Nash equilibria exist when the extended utility function U_i is concave in X_i . However, Observation 1 indicates that a Nash equilibrium may fail to exist when U_i is strictly convex in X_i . Indeed, take a game, $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ with no pure strategy Nash equilibrium, like Matching Pennies, and consider its mixed extension $\Gamma = \langle N, (X_i)_{i \in N}, (U_i)_{i \in N} \rangle$, where for all players, their extended utility function is strictly convex. Then, for any player $i \in N$ and for any profile of mixed strategies x_{-i} of the other players, the set of i 's best responses $\mathcal{B}_i(x_{-i})$ consists of only degenerate mixed strategies. Since G has no pure strategy Nash equilibrium, we conclude that Γ does not have a Nash equilibrium.

Observation 2 It is also important to note that, unlike in the standard expected utility case, a player's mixed strategy x_i^* may very well be a best response to some profile x_{-i}^* of the other players' mixed strategies and at the same time may assign positive probability to an action that (when regarded as a degenerate mixed strategy) is not a best response to x_{-i}^* . Formally, it may very well be the case that

$$U_i((x_k^*)_{k \in N}) \geq U_i((x_i, x_{-i}^*)) \quad \text{for all } x_i \in X_i$$

and yet

$$U_i((a_i, x_{-i}^*)) < U_i((x_k^*)_{k \in N}) \\ \text{for some } a_i \text{ such } x_i^*(a_i) > 0.$$

This will typically occur when the function U_i is strictly concave in X_i .

The definition mixed strategy equilibrium requires from each strategy in the equilibrium profile that it be a best response to the other strategies. Corollary 1 stated that when preferences have the expected utility form, each mixed strategy in a mixed strategy equilibrium is also a probability mixture over best responses to the other strategies in the profile. This result allowed us to interpret a mixed strategy Nash equilibrium as a profile of beliefs, rather than as a profile of probability mixtures. As explained in Observation 2, however, when preferences over mixed strategies are not expected utility preferences, a mixture over best responses is not necessarily a best response. Therefore, Corollary 1 does not extend to the mixed extension where preferences are not of the expected utility form.

In this setup, however, one can still interpret a player's mixed strategy as a belief entertained by the other players about the actions chosen by that player. And a profile of such beliefs will be in equilibrium if the probability distribution over the player's actions that represents i 's beliefs is obtained as a mixture of best responses of this player to his beliefs about the other players' actions. With this idea in mind, Crawford (1990)

defined the notion of an equilibrium in beliefs. Before we formally present his definition we need to introduce some notation.

Since when the extended utility functions U_i are concave in i 's own strategy a best response to a given profile of the other players' strategies may be a non-degenerate mixed strategy, a mixture of best responses will typically be a mixture over non-degenerate mixed strategies. This mixture induces a probability distribution over actions in a natural way by reducing the compound mixture to a simple mixture. This induced probability distribution can be interpreted as a belief over the actions ultimately chosen. For example, in Matching Pennies, if player 1 believes that there is a probability of 1/2 that player 2 will choose the mixed strategy (1/3, 2/3) and a probability of 1/2 that player 2 will choose the mixed strategy (2/3, 1/3), then player 1 believes that player 2 will choose each one of his two actions with equal probability. More generally, if player i assigns probability p_k to the event that player j will choose mixed strategy $x^k \in X_j$, for $k = 1, \dots, K$, then player i 's beliefs about player j 's actions are given by $\sum_{k=1}^K p_k x^k \in X_j$. That is, for each action $a_j \in A_j$ of player j , player i believes that player j will choose a_j with probability $\sum_{k=1}^K p_k x^k(a_j)$. For each set $T \subset X_i$ of mixed strategies, let $D[T] \subset X_i$ denote the set of probability distributions over i 's actions that are induced by mixtures over elements of T . With this notation in hand, we can define the concept of equilibrium in beliefs.

Definition 5 Let $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ be a game. For each $i \in N$, let $\mathcal{B}_i : X \rightarrow X_i$ be the best response correspondence in the mixed extension of G . The profile of beliefs $(x_k^*)_{k \in N} \in \times_{k \in N} \Delta(A_k)$ is an *equilibrium in beliefs* if

$$x_i^* \in D[\mathcal{B}_i(x_{-i}^*)] \quad \text{for all } i \in N.$$

An equilibrium in beliefs is a profile of beliefs $(x_k^*)_{k \in N}$. For each $i \in N$, x_i^* is the common belief of the players other than i about player i 's choice of actions. In order for this profile of beliefs to be in equilibrium, we require that for each player $i \in N$ all the other players believe that i chooses a mixed strategy that is a best response to his

beliefs, which are given by $(x_k^*)_{k \in N \setminus \{i\}}$, about the other players' choices actions. In other words, x_i^* must be a convex combination of best responses of i to $(x_k^*)_{k \in N \setminus \{i\}}$.

Example 2 Consider again the mixed extension of Matching Pennies $\langle N, (X_i)_{i \in N}, (U_i)_{i \in N} \rangle$ where the set of players is $N = \{1, 2\}$, the sets of mixed strategies are $X_1 = \{(p_H, p_T) \geq (0, 0) : p_H + p_T = 1\}$ and $X_2 = \{(q_H, q_T) \geq (0, 0) : q_H + q_T = 1\}$, and the utility functions are now given by $U_1((p_H, p_T), (q_H, q_T)) = (p_H q_H)^2 + (p_T q_T)^2 - p_H q_T - p_T q_H$ and $U_2((p_H, p_T), (q_H, q_T)) = (p_H q_T)^2 + (p_T q_H)^2 - p_H q_H - p_T q_T$. Since the utility functions are strictly convex in the players' own mixed strategies, the best response to any strategy of the opponent is a pure strategy. In particular, one can verify that

$$\mathcal{B}_1(q_H, q_T) = \begin{cases} (0, 1) & \text{if } q_H > q_T \\ \{(1, 0), (0, 1)\} & \text{if } q_H = q_T \\ (1, 0) & \text{if } q_H < q_T \end{cases}$$

and

$$\mathcal{B}_1(p_H, p_T) = \begin{cases} (0, 1) & \text{if } p_H > p_T \\ \{(1, 0) \cdot (0, 1)\} & \text{if } p_H = p_T \\ (1, 0) & \text{if } p_H < p_T. \end{cases}$$

It can also be verified that $((p_H^*, p_T^*), (q_H^*, q_T^*)) = ((1/2, 1/2), (1/2, 1/2))$ is an equilibrium in beliefs. Indeed, for both $i = 1, 2$, $(1/2, 1/2) \in X_i$ is a convex combination of $(1, 0)$ and $(0, 1)$, which are both in $\mathcal{B}_j(1/2, 1/2), j \neq i$. In this equilibrium,

1. Player 1 believes that player 2 will choose $(1, 0)$ with probability 1/2, and $(0, 1)$ with probability 1/2.
2. Therefore player 1 believes that player 2 will ultimately choose H and T , each with probability 1/2.
3. Given these beliefs, player 1's only best replies are $(1, 0)$ and $(0, 1)$, and
4. Player 2 believes that player 1 will choose each one with probability 1/2. As a result,

5. Player 2 believes that player 1 will ultimately choose H and T each with probability 1/2.
6. Given these beliefs, player 2's only best replies are $(1, 0)$ and $(0, 1)$, and
1. Player 1 believes that player 2 will choose $(1, 0)$ with probability 1/2, and $(0, 1)$ with probability 1/2.

The following result is a direct implication of the definition of an equilibrium in beliefs.

Proposition 2 (Crawford 1990)) Let $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ be a strategic game, and $\Gamma = \langle N, (X_i)_{i \in N}, (U_i)_{i \in N} \rangle$ be the mixed extension of G , where U_i is continuous but not necessarily multilinear.

1. Every mixed strategy Nash equilibrium of G is an equilibrium in beliefs.
2. If for all $i \in N$, U_i is quasiconcave in X_i , then every equilibrium in beliefs is a mixed strategy Nash equilibrium of G .

Proof

1. Since $\mathcal{B}_i(x_{-i}^*) \subset D[\mathcal{B}_i(x_{-i}^*)]$ for all $i \in N$, every Nash equilibrium is an equilibrium in beliefs.
2. When the utility function U_i is quasiconcave in i 's mixed strategy, the set of best responses $\mathcal{B}_i(x_{-i}^*)$ is a convex set. Therefore, $D[\mathcal{B}_i(x_{-i}^*)] = \mathcal{B}_i(x_{-i}^*)$, and any equilibrium in beliefs is a Nash equilibrium. $\square$

Crawford (1990) shows that although some games have no Nash equilibrium, every game has an equilibrium in beliefs.

Correlated Equilibrium

In the mixed extension of a game, players do not choose their actions directly, but rather choose probability distributions over their action sets according to which the actions are ultimately selected. The important feature about these probability distributions is that they represent independent random variables. The realization of one player's random variable does not give any information about the

realization of the other players' random variables. There is nothing in the bare notion of equilibrium, however, that requires players' behavior to be independent. The basic feature of an equilibrium is that each player is best responding to the behavior of others, and that each player is free to choose any action in his action set. But one thing is that players can, if they so wish, change their behavior without the consent of others, and another different thing is to expect players' choices to be independent. Therefore, one could ask what would happen if the random devices players use to ultimately choose their actions were correlated. In that case, knowledge of the realization of one's random device would provide some partial information about the realization of the other players' random devices, and therefore of their choices. In equilibrium, a player should take this information into account. To illustrate this point, consider the game of Chicken.

		Driver 2	
		Slow Down	Speed up
Driver 1	Slow Down	6, 6	2, 7
	Speed up	7, 2	0, 0

This game has two pure action Nash equilibria, and one equilibrium in mixed strategies. According to the mixed strategy Nash equilibrium, each player chooses Slow Down with probability 2/3 and Speed Up with probability 1/3. This mixed strategy equilibrium can be implemented by the following random device. Consider two random variables S_1 and S_2 , whose joint distribution is given by the following table:

Driver 1 chooses his action as a function of the realization of S_1 and Driver 2 chooses his action as a function of the realization of S_2 . (Neither player is informed of the realization of the other player's random variable.) In particular, Driver 1 chooses Slow Down if $S_1 = 1$ and Speed Up otherwise. Similarly, Driver 2 chooses Slow Down if $S_2 = 1$, and Speed Up otherwise. Note that according to this pattern behavior, each player chooses to slow down with probability 2/3. But more importantly, since S_1 and S_2 are independent random variables, knowledge of the realization of one random variable does not give any information about the realization of the other one. Therefore, after

Static Games, Table 1 A random device

	S_2	
	1	2
S_1	1	4/9 2/9
	2	2/9 1/9

Driver 1 learns the realization of S_1 , he still believes that Driver 2 will choose Slow Down with probability 2/3 and consequently any choice is optimal, in particular the one described above. Similarly, after Driver 2 learns the realization of S_2 , he still believes that Driver 1 will choose to slow down with probability 2/3, and his planned behavior continues to be optimal.

But what would happen if the joint distribution of S_1 and S_2 , was not as presented in Table 1, but rather as follows?

		S_2	
		1	2
S_1	1	1/3	1/3
	2	1/3	0

To answer this question, assume that both players still choose their actions according to the previous pattern of behavior: Driver 1 chooses Slow Down if $S_1 = 1$, and Speed Up otherwise. The same holds for Driver 2. As a result, it is still true that each player chooses Slow Down with probability 2/3 and Speed Up with probability 1/3. However, since this time the conditioning random variables S_1 and S_2 are not independent, knowledge of the realization of S_1 affects the beliefs of Driver 1 about the probability with which Driver 2 chooses his actions. In particular, if $S_1 = 1$, Driver 1 updates his beliefs and assigns probability 1/2 to Driver 2 choosing either action, and consequently, Driver 1's only optimal action is Slow Down, which is precisely the choice dictated by the above pattern of behavior. Similarly, if $S_1 = 2$, Driver 1 should update his beliefs and assign probability one that Driver 2 will choose Slow Down. Consequently, Driver 1's best reply is to follow the above pattern of behavior and choose Speed Up. One can see that, given that the players know that the random variables S_1 and S_2 are correlated and they use this information accordingly, there is no incentive for either of them to deviate from the proposed pattern of behavior. Therefore, we can say that this pattern of behavior is an

equilibrium. This notion of a correlated equilibrium was introduced in Aumann (1974). Before we give a formal definition we introduce the concept of a correlated strategy profile, which will play a central role not only in this section, but in the next one as well.

Definition 6 Let $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ be a game. A *correlated strategy profile* in G consists of

- A finite probability space (Ω, π)
- For each player $i \in N$, a partition $\mathcal{P}_i$ of Ω into events of positive probability
- For each player $i \in N$, a function $\sigma_i : \Omega \rightarrow A_i$ which is measurable with respect to $\mathcal{P}_i$.

A correlated strategy profile is a description of what players do and know while playing the game G . The collection $\langle (\Omega, \pi), (\mathcal{P}_i)_{i \in N} \rangle$ represents the random devices used by the players to ultimately choose their actions. The underlying probability space that governs the players' random devices is (Ω, π) . Ω is the set of states, and for each state ω , $\pi(\omega)$ is the probability that ω occurs. For each $i \in N$, the partition $\mathcal{P}_i$ represents player i 's information. Each element of the partition represents a different realization of the random device used by i to choose his action. Two states that belong to the same element of the partition $\mathcal{P}_i$ cannot be distinguished by i , while two states that belong to different partition cells can be distinguished by him. For each player i , $\sigma_i : \Omega \rightarrow A_i$ is the random variable that describes player i 's choice of action, $\sigma_i(\omega)$ being the action chosen by him at state ω . The measurability of σ_i with respect to $\mathcal{P}_i$ formalizes the requirement that the actions chosen by player i depend only on his information about the state of the world. Therefore, for any two states that belong to the same element of his partition, the actions chosen by i at those states must be the same. That is, for any $\omega, \omega' \in P \in \mathcal{P}_i$ we have $\sigma_i(\omega) = \sigma_i(\omega')$.

For example, the correlated strategy profile described earlier for the game of chicken can be formalized as $\langle (\Omega, \pi), (\mathcal{P}_i)_{i \in N}, (\sigma_i)_{i \in N} \rangle$, where $N = \{I, II\}$, and

- $\Omega = \{(1, 1), (1, 2), (2, 1)\}$
- $\pi(\omega) = 1/3$ for all $\omega \in \Omega$
- $\mathcal{P}_I = \{\{(1, 1), (1, 2)\}, \{(2, 1)\}\}$ and $\mathcal{P}_{II} = \{\{(1, 1), (2, 1)\}, \{(1, 2)\}\}$
- $\sigma_I(\omega) = \begin{cases} \text{Slow Down} & \text{if } \omega \in \{(1, 1), (1, 2)\} \\ \text{Speed up} & \text{if } \omega \in \{(2, 1)\} \end{cases}$
- $\sigma_{II}(\omega) = \begin{cases} \text{Slow Down} & \text{if } \omega \in \{(1, 1), (2, 1)\} \\ \text{Speed up} & \text{if } \omega \in \{(1, 2)\} \end{cases}$

According to this correlated strategy profile, there are three equally likely states, and the players can distinguish only one component of the state, namely the realization of their random variable. The players' actions are described by the functions σ_I and σ_{II} which depend only on the respective player's information.

In what follows we denote by $\sigma : \Omega \rightarrow A$ the function that associates with each $\omega \in \Omega$ the action profile induced by the strategies σ_k , for $k \in N$. That is, $\sigma = (\sigma_k)_{k \in N}$. Also, for any $i \in N$, $\sigma_{-i} = (\sigma_k)_{k \in N \setminus \{i\}}$ so that $\sigma = (\sigma_{-i}, \sigma_i)$. We are interested in correlated strategy profiles in which no player benefits by altering his behavior. These special profiles are introduced in the following definition.

Definition 7 Let $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ be a strategic game. A *correlated equilibrium* of G is a correlated strategy $\langle (\Omega, \pi), (\mathcal{P}_i)_{i \in N}, (\sigma_i)_{i \in N} \rangle$ such that for every $i \in N$ and every function $\tau_i : \Omega \rightarrow A_i$ that is measurable with respect to $\mathcal{P}_i$,

$$\sum_{\omega \in \Omega} \pi(\omega) u_i(\sigma_{-i}(\omega), \sigma_i(\omega)) \geq \sum_{\omega \in \Omega} \pi(\omega) u_i(\sigma_{-i}(\omega), \tau_i(\omega)). \quad (6)$$

The value $v_i = \sum_{\omega \in \Omega} \pi(\omega) u_i(\sigma_{-i}(\omega), \sigma_i(\omega))$ is player i 's correlated equilibrium payoff.

In a correlated strategy profile each player plans to condition his choice of action on the realization of a random variable, and the players' random variables may be correlated. A correlated strategy profile is a correlated equilibrium if no player can find an alternative way to condition his choice on the same random device, so that his

expected utility is increased. Note that the player presumably chooses his strategy (his way to condition his actions on the outcomes of the random device) before he learns the realization of the device. Nonetheless, he evaluates the outcomes generated by the players' strategies by taking into account the precise correlation of the random devices on which outcomes players are conditioning their behavior. Although strictly speaking mixed strategy Nash equilibria are not correlated equilibria, they do induce a correlated equilibrium distribution over action profiles. In order to state this claim, we need the following definition.

Definition 8 Let $\langle (\Omega, \pi), (\mathcal{P}_i, \sigma_i)_{i \in N} \rangle$ be a correlated strategy profile for G . Its induced probability distribution over action profiles is given by the function $p : A \rightarrow [0, 1]$ defined by

$$\begin{aligned} p(a) &= \pi(\{\omega \in \Omega : \sigma(\omega) = a\}) \\ &= \sum_{\{\omega \in \Omega : \sigma(\omega) = a\}} \pi(\omega) \quad \text{for all } a \in A \end{aligned}$$

Proposition 3 Let $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ be a strategic game, and let $x = (x_1, \dots, x_n)$ be a mixed strategy Nash equilibrium of G . Then, there is a correlated equilibrium $\langle (\Omega, \pi), (\mathcal{P}_i)_{i \in N}, (\sigma_i)_{i \in N} \rangle$, whose induced probability distribution over action profiles is the same as x 's distribution.

Proof Let $\langle (\Omega, \pi), (\mathcal{P}_i)_{i \in N}, (\sigma_i)_{i \in N} \rangle$ be defined as follows:

$$\begin{aligned} \Omega &= A \\ \pi(a) &= \prod_{i \in N} x_i(a_i) \\ \mathcal{P}_i(a) &= \{b \in A : b_i = a_i\} \\ \sigma_i(a) &= a_i. \end{aligned}$$

We claim that $\langle (\Omega, \pi), (\mathcal{P}_i)_{i \in N}, (\sigma_i)_{i \in N} \rangle$ is a correlated equilibrium whose probability distribution is the same as x 's distribution. Let $i \in N$. Since x is a mixed strategy Nash equilibrium, we know by Lemma 1 that for all $a_i \in A_i$

$$\begin{aligned} \text{if } x_i(a_i) > 0 \text{ then } U_i(x) &= U_i(a_i, x_{-i}) \\ \text{if } x_i(a_i) = 0 \text{ then } U_i(x) &\geq U_i(a_i, x_{-i}). \end{aligned}$$

Consequently, for all $a_i \in A_i$

$$x_i(a_i)U_i(a_i, x_{-i}) \geq x_i(a_i)U_i(b_i, x_{-i}) \quad \text{for all } b_i \in A_i. \quad (7)$$

Now let $\tau_i : A \rightarrow A_i$ be a function that is measurable with respect to $\mathcal{P}_i$. Let $a_{-i} \in A_{-i}$ be a fixed profile of actions for players other than i . Letting $b_i = \tau_i(a_i, a_{-i})$, Eq. (7) implies that

$$\begin{aligned} x_i(a_i)U_i(a_i, x_{-i}) &\geq x_i(a_i)U_i(\tau_i(a_i, a_{-i}), x_{-i}) \\ &\quad \text{for all } a_i \in A_i \end{aligned}$$

Adding over all $a_i \in A_i$,

$$\begin{aligned} \sum_{a_i \in A_i} x_i(a_i)U_i(a_i, x_{-i}) &\geq \\ \sum_{a_i \in A_i} x_i(a_i)U_i(\tau_i(a_i, a_{-i}), x_{-i}). \end{aligned}$$

Taking into account the definition of $U_i(a_i, x_{-i})$ and $U_i(\tau_i(a_i, x_{-i}))$, and using the measurability of τ_i with respect to $\mathcal{P}_i$, we get

$$\begin{aligned} \sum_{a_i \in A_i} x_i(a_i) \sum_{a_{-i} \in A_{-i}} \left(\prod_{j \in N \setminus \{i\}} x_j(a_j) \right) u_i(a_i, a_{-i}) \\ &\geq \sum_{a_i \in A_i} x_i(a_i) \sum_{a_{-i} \in A_{-i}} \left(\prod_{j \in N \setminus \{i\}} x_j(a_j) \right) u_i(\tau_i(a_i, a_{-i}), a_{-i}) \\ &= \sum_{a \in A} \left(\prod_{j \in N} x_j(a_j) \right) u_i(a_i, a_{-i}) \\ &\geq \sum_{a \in A} \left(\prod_{j \in N} x_j(a_j) \right) u_i(\tau_i(a), a_{-i}) \\ &= \sum_{a \in A} \pi(a) u_i(a_i, a_{-i}) \geq \sum_{a \in A} \pi(a) u_i(\tau_i(a), a_{-i}) \\ &= \sum_{a \in A} \pi(a) u_i(\sigma_i(a), \sigma_{-i}(a)) \geq \sum_{a \in A} \pi(a) u_i(\tau_{-i}(a), \sigma_{-i}(a)). \end{aligned}$$

This shows that $\langle (\Omega, \pi), (\mathcal{P}_i)_{i \in N}, (\sigma_i)_{i \in N} \rangle$ is a correlated equilibrium of G . Its induced probability distribution over action profiles is

$$\begin{aligned}
p(a) &= \pi(\{b \in A : \sigma(b) = a\}) \\
&= \pi(\{b \in A : b = a\}) \\
&= \pi(a) \\
&= \prod_{i \in N} x_i(a_i).
\end{aligned}$$

□

Although a correlated strategy profile consists of a randomizing device used by the players, it turns out that the only feature of the device that determines whether or not the correlated strategy profile constitutes a correlated equilibrium is its induced probability distribution over the action profiles. This is shown by the next proposition.

Proposition 4 *Let $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ be a finite strategic game. Every correlated equilibrium probability distribution over action profiles can be obtained in a correlated equilibrium of G in which*

- $\Omega = A$
- $\mathcal{P}_i(a) = \{b \in A : b_i = a_i\}$.

Proof Let $\langle (\Omega', \pi') (\mathcal{P}'_i, \sigma'_i)_{i \in N} \rangle$ be a correlated equilibrium of G . Consider the correlated strategy profile $\langle (\Omega, \pi), (\mathcal{P}_i, \sigma_i)_{i \in N} \rangle$ defined by

- $\Omega = A$
- $\pi(a) = \pi'(\{\omega \in \Omega' : \sigma'(\omega) = a\})$ for each $a \in A$
- $\mathcal{P}_i(a) = \{b \in A : b_i = a_i\}$ for each $i \in N$ and for each $a \in A$
- $\sigma_i(a) = a_i$ for each $i \in N$.

It is clear that this correlated strategy profile induces the required distribution over action profiles. Indeed,

$$\begin{aligned}
p(a) &= \pi(\{\omega \in \Omega : \sigma(\omega) = a\}) \\
&= \pi(\{a' \in A : a' = a\}) \\
&= \pi(a) \\
&= \pi'(\{\omega \in \Omega' : \sigma'(\omega) = a\}).
\end{aligned}$$

It remains to show that this profile is a correlated equilibrium. Take a function $\tau_i : A \rightarrow A_i$ that is measurable with respect to $\mathcal{P}_i$. Define $\tau'_i : \Omega' \rightarrow A_i$ by $\tau'_i(\omega) = \tau_i(\sigma'(\omega)) = \tau_i(\sigma'_{-i}(\omega), \sigma'_i(\omega))$. The function τ'_i is measurable with respect to $\mathcal{P}'_i$. Indeed,

if $\omega' \in P'_i(\omega)$, then $\sigma'_i(\omega') = \sigma'_i(\omega)$ by measurability of σ'_i with respect to $\mathcal{P}'_i$. Therefore, by definition of $\mathcal{P}_i$, $\mathcal{P}_i(\sigma'(\omega')) = \mathcal{P}_i(\sigma'(\omega))$, and both $\sigma'_i(\omega')$ and $\sigma'_i(\omega)$ belong to the same element of $\mathcal{P}_i$. Since τ_i is measurable with respect to $\mathcal{P}_i$, we conclude that $\tau'_i(\omega') = \tau_i(\sigma'(\omega')) = \tau_i(\sigma'(\omega)) = \tau'_i(\omega)$.

Also,

$$\begin{aligned}
&\sum_{\omega \in \Omega} \pi(\omega) u_i(\sigma_{-i}(\omega), \tau_i(\omega)) \\
&= \sum_{a \in A} \pi(a) u_i(a_{-i}, \tau_i(a)) \\
&= \sum_{a \in A} \sum_{\{\omega \in \Omega' : \sigma'(\omega) = a\}} \pi'(\omega) u_i(\sigma'_{-i}(\omega), \tau_i(\sigma'(\omega))) \\
&= \sum_{a \in A} \sum_{\{\omega \in \Omega' : \sigma'(\omega) = a\}} \pi'(\omega) u_i(\sigma'_{-i}(\omega), \tau'_i(\omega)) \\
&= \sum_{\omega \in \Omega'} \pi'(\omega) u_i(\sigma'_{-i}(\omega), \tau'_i(\omega))
\end{aligned}$$

In particular, for $\tau_i = \sigma_i$,

$$\begin{aligned}
&\sum_{\omega \in \Omega} \pi(\omega) u_i(\sigma_{-i}(\omega), \sigma_i(\omega)) \\
&= \sum_{\omega \in \Omega'} \pi'(\omega) u_i(\sigma'_{-i}(\omega), \sigma'_i(\omega)).
\end{aligned}$$

Since $\langle (\Omega', \pi'), (\mathcal{P}'_i, \sigma'_i)_{i \in N} \rangle$ is a correlated equilibrium,

$$\begin{aligned}
&\sum_{\omega \in \Omega'} \pi'(\omega) u_i(\sigma'_{-i}(\omega), \sigma'_i(\omega)) \\
&\geq \sum_{\omega \in \Omega'} \pi'(\omega) u_i(\sigma'_{-i}(\omega), \tau'_i(\omega))
\end{aligned}$$

and therefore

$$\begin{aligned}
&\sum_{\omega \in \Omega} \pi(\omega) u_i(\sigma_{-i}(\omega), \sigma_i(\omega)) \\
&\geq \sum_{\omega \in \Omega} \pi(\omega) u_i(\sigma_{-i}(\omega), \tau_i(\omega)).
\end{aligned}$$

□

Rationality, Correlated Equilibrium and Equilibrium in Beliefs

As mentioned earlier, Nash equilibrium and correlated equilibrium are two examples of what is known as solution concepts. Solution concepts assign to each game a pattern behavior for the players in the game. The interpretation of these

patterns of behavior is not always explicit, but it is fair to say that they are usually interpreted either as descriptions what rational people do, or as prescriptions of what rational people should do. There is a growing literature that tries to connect various game theoretic solution concepts to the idea of rationality. Rationality is generally understood as the characteristic of a player who chooses an action that maximizes his preferences, given his information about the environment in which he acts. Part of the information a player has is represented by his beliefs about the behavior of other players, their beliefs about the behavior of other players, and so on. So when one speaks of the rationality players, one needs to take into account their *epistemic state*. There is a formal framework which is appropriate for discussing the actions, knowledge, beliefs and rationality players. Namely, the framework of a correlated strategy profile. As defined in Sect. “Correlated Equilibrium”, a correlated strategy profile in a game G consists of

- A finite probability space (Ω, π)
- For each player $i \in N$ a partition $\mathcal{P}_i$ of Ω into events of positive probability
- For each player $i \in N$ a function $\sigma_i : \Omega \rightarrow A_i$ which is measurable with respect to $\mathcal{P}_i$.

For the present discussion we interpret a correlated strategy profile $\langle (\Omega, \pi), (\mathcal{P}_i)_{i \in N}, (\sigma_i)_{i \in N} \rangle$ as a description of the players' behavior and beliefs, as observed by an outside observer. The set Ω is the set of possible states of the world and π is the prior probability on Ω shared by all the players. For each player $i \in N$, $\mathcal{P}_i$ is a partition of Ω that represents i 's information. At state $\omega \in \Omega$, player i is informed not of the state that actually occurred, but of the element $\mathcal{P}_i(\omega)$ of his partition that contains ω . Player i then uses this information and his prior π to update his beliefs about the true state of the world. Finally, the function σ_i represents the actions taken by player i at each state. In particular, $\sigma_i(\omega)$ is the action chosen by i at state ω . Although a correlated equilibrium can be interpreted as a correlated strategy profile *prescribed* by a given solution concept (that of a correlated equilibrium), here we want to interpret a correlated strategy profile as a description of what players actually do and

believe. Although players cannot freely choose their beliefs (in the same way as they cannot choose their preferences), they can choose their actions. Furthermore, they have no obligation to behave according to the specified correlated strategy profile. However, ultimately players do behave in a certain way and that behavior is what is represented by the given correlated strategy profile.

Once we fix a correlated strategy profile we can address the rationality of the players. Formally,

Definition 9 Player $i \in N$ is *Bayes rational* at $\omega \in \Omega$ if his expected payoff at ω , $E(u_i(\sigma) | \mathcal{P}_i)(\omega)$, is at least as large as the amount $E(u_i(\sigma_{-i}, a_i) | \mathcal{P}_i)(\omega)$ that he would have got had he chosen action $a_i \in A_i$ instead of $\sigma_i(\omega)$.

In other words, player i is rational at a given state of the world if the action $\sigma_i(\omega)$ he chooses at that state maximizes his expected utility given his information, $\mathcal{P}_i(\omega)$, and, in particular, given his beliefs about the actions of the other players.

As before, for any finite set T , let $\Delta(T)$ be the set of all probability distributions on T . The beliefs of player i about the actions of the other players are represented by his conjectures. A conjecture of i is a probability distribution $\psi_i \in \Delta(A_{-i})$ over the elements of A_{-i} . For any $j \neq i$, the marginal of ψ_i on A_j is the conjecture of i about j induced by ψ_i . Given a correlated strategy profile $\langle (\Omega, \pi), (\mathcal{P}_i)_{i \in N}, (\sigma_i)_{i \in N} \rangle$, one can determine the conjectures that each player is entertaining at each state of the world about the actions of the other players. These conjectures are given by the following definition.

Definition 10 Given a correlated strategy profile $\langle (\Omega, \pi), (\mathcal{P}_i, \sigma_i)_{i \in N} \rangle$, the conjectures of $i \in N$ about the other players' actions are given by the function $\phi_i : \Omega \rightarrow \Delta(A_{-i})$ defined by

$$\phi_i(\omega)(a_{-i}) = \frac{\pi[\{\omega' \in \mathcal{P}_i(\omega) : \sigma_{-i}(\omega') = a_{-i}\}]}{\pi[\mathcal{P}_i(\omega)]}.$$

For each ω , $\phi_i(\omega) \in \Delta(A_{-i})$ is the conjecture of i at ω . For $j \neq i$, the marginal of $\phi_i(\omega)$ on A_j is the conjecture of i at ω about j 's actions.

Given a correlated strategy profile, we can speak about what each player knows. The object of knowledge are called *events*, which are the subsets of the set of states of the world Ω . We say that player i knows event $E \subset \Omega$ at state ω , if $P_i(\omega) \subset E$. That is, i knows E at ω if whatever state he deems possible at ω is in E .

The next result, proved by Aumann and Brandenburger (1995), shows a remarkable relationship between the rationality players and the concept of equilibrium in beliefs.

Theorem 2 Fix a two-person game, $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$, and let $\langle (\Omega, \pi), (\mathcal{P}_i)_{i \in N}, (\sigma_i)_{i \in N} \rangle$ be a correlated strategy profile for G . Let $\psi_1 \in \Delta(A_1)$ and $\psi_2 \in \Delta(A_2)$ be two conjectures, one about player 1's actions and the other about player 2's actions. Assume that at some state $\omega \in \Omega$ each player knows that the other is rational and that their conjectures at ω are $(\phi_1(\omega), \phi_2(\omega)) = (\psi_2, \psi_1)$. Then, (ψ_1, ψ_2) is an equilibrium in beliefs.

Proof The fact that player i knows at ω that j 's conjecture is ψ_j means that

$$\mathcal{P}_i(\omega) \subset \{\omega' \in \Omega : \phi_j(\omega')(a_i) = \psi_j(a_i) \text{ for all } a_i \in A_i\}.$$

Therefore

$$\phi_j(\omega)(a_i) = \psi_j(a_i) \text{ for all } a_i \in A_i. \quad (8)$$

Given Proposition 2 and Corollary 1, we need to show that if $\psi_i(a_i^*) > 0$, a_i^* is a best response to ψ_j , for $i, j = 1, 2, i \neq j$. For this purpose, assume that $\psi_i(a_i^*) > 0$ for some $a_i^* \in A_i$. Then, by definition of ϕ_j and (8), $\phi_j(\omega)(a_i^*) = \pi[\{\omega' \in \mathcal{P}_j(\omega) : \sigma_i(\omega') = a_i^*\}] > 0$. Consequently, there is $\omega' \in \mathcal{P}_j(\omega)$ such that $\sigma_i(\omega') = a_i^*$. Since player knows at ω that player i is rational,

$$\begin{aligned} \omega' \in \mathcal{P}_j(\omega) &\subset \{\omega'' \in \Omega : E[u_i(\sigma) | \mathcal{P}_i](\omega'') \\ &\geq E[u_i(\sigma_{-i}, a_i) | \mathcal{P}_i](\omega'') \text{ for all } a_i \in A_i\}. \end{aligned}$$

Therefore,

$$E[u_i(\sigma) | \mathcal{P}_i](\omega') \geq E[u_i(\sigma_{-i}, a_i) | \mathcal{P}_i](\omega') \text{ for all } a_i \in A_i$$

and since $\sigma_i : \Omega \rightarrow A_i$ is measurable with respect to $\mathcal{P}_i$, $\sigma_i(\omega') = a_i^*$ is the action that player i chooses at all states in $\mathcal{P}_i(\omega')$. Then we can write

$$E[u_i(\sigma_{-i}, a_i^*) | \mathcal{P}_i](\omega') \geq E[u_i(\sigma_{-i}, a_i) | \mathcal{P}_i](\omega') \text{ for all } a_i \in A_i.$$

That is, for all $a_i \in A_i$

$$\begin{aligned} &\sum_{\omega'' \in \mathcal{P}_j(\omega')} \frac{\pi(\omega'')}{\pi(\mathcal{P}_i(\omega'))} u_i(\sigma_{-i}(\omega''), a_i^*) \\ &\geq \sum_{\omega'' \in \mathcal{P}_j(\omega')} \frac{\pi(\omega'')}{\pi(\mathcal{P}_i(\omega'))} u_i(\sigma_{-i}(\omega''), a_i^*) \\ &\sum_{a_j \in A_j} \sum_{\omega'' \in \mathcal{P}_j(\omega') \atop \sigma_j(\omega'') = a_j} \frac{\pi(\omega'')}{\pi(\mathcal{P}_i(\omega'))} u_i(a_j, a_i^*) \\ &\geq \sum_{a_j \in A_j} \sum_{\omega'' \in \mathcal{P}_j(\omega') \atop \sigma_j(\omega'') = a_j} \frac{\pi(\omega'')}{\pi(\mathcal{P}_i(\omega'))} u_i(a_j, a_i) \\ &\sum_{a_j \in A_j} \frac{\pi[\{\omega'' \in \mathcal{P}_i(\omega') : \sigma_j(\omega'') = a_j\}]}{\pi(\mathcal{P}_i(\omega'))} u_i(a_j, a_i^*) \\ &\geq \sum_{a_j \in A_j} \frac{\pi[\{\omega'' \in \mathcal{P}_i(\omega') : \sigma_j(\omega'') = a_j\}]}{\pi(\mathcal{P}_i(\omega'))} u_i(a_j, a_i) \\ &\sum_{a_j \in A_j} \phi_i(\omega')(a_j) u_i(a_j, a_i^*) \\ &\geq \sum_{a_j \in A_j} \phi_i(\omega')(a_j) u_i(a_j, a_i) \end{aligned} \quad (9)$$

Since $\omega' \in \mathcal{P}_i(\omega)$ and player j knows at ω that i 's conjecture is ψ_j , then

$$\omega' \in \mathcal{P}_i(\omega) \subset \{\omega'' \in \Omega : \phi_i(\omega'')(a_j) = \psi_j(a_j) \text{ for all } a_j \in A_j\}.$$

Therefore $\phi_i(\omega')(a_j) = \psi_j(a_j)$ for all $a_j \in A_j$, or $\phi_i(\omega') = \psi_j$. That is, i 's conjecture at ω' about j 's actions is ψ_j . Consequently, substituting in (9),

$$\sum_{a_j \in A_j} \psi_j(a_j) u_i(a_j, a_i^*) \geq \sum_{a_j \in A_j} \psi_j(a_j) u_i(a_j, a_i) \quad \forall a_i \in A_i.$$

That is, a_i^* is a best response to player i 's beliefs about j 's actions. $\square$

The only assumptions required by Theorem 2 is that players know they are rational, and that they know each other's conjectures. In a correlated strategy profile for a two-player game, there is only one player entertaining a conjecture about the actions of player 1, namely, player 2. Similarly, player 1 is the only one who entertains a conjecture about the actions of player 2. In an n -person game, with $n > 2$, for each player, there is more than one player entertaining a conjecture about his actions. Therefore, since an equilibrium in beliefs consists of a profile of beliefs, each of which is shared by $n - 1$ players, a generalization of Theorem 2 would require the players' beliefs about player i 's actions, for $i \in N$, to be identical. In order to obtain these common beliefs it is not sufficient to assume that players know each other's conjectures. One need to strengthen this assumption. Also, in an equilibrium in beliefs, the common belief about player i 's actions assigns positive probability only to best responses to i 's conjectures about the choices of the other players. Furthermore, i 's conjectures about the other players' choices is the product of his beliefs about each of the other players. In other words, an equilibrium in beliefs implicitly assumes that players believe that the other players' choices are independent. Aumann and Brandenburger (1995) show that one way to obtain common conjectures and, simultaneously, that players believe that the other players act independently, is to assume that players' conjectures are commonly known. This surprising and deep result is stated in the next theorem.

Theorem 3 *Let $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ be a strategic game, and let $\langle (\Omega, \pi), (\mathcal{P}_i)_{i \in N}, (\sigma_i)_{i \in N} \rangle$ be a correlated strategy profile for G . Also let $(\psi_i)_{i \in N} \in \times_{i \in N} \Delta(A_{-i})$ be a profile of conjectures, one for each player. Assume that at some state $\omega \in \Omega$ each player knows that the others are rational. Further, assume that at ω their conjectures are commonly known to be $(\psi_i)_{i \in N}$. Then, for each j , all the conjectures ψ_j of players j other than j , induce the same belief $\varphi_j \in \Delta(A_j)$ about j 's actions, and the resulting profile of beliefs, $(\varphi_i)_{i \in N}$, is an equilibrium in beliefs.*

Rationality and Correlated Equilibrium

The previous result shows a surprising relationship between the players' rationality and the concept of equilibrium in beliefs. If at some state of the world players know that everybody is rational, and their conjectures are commonly known at that state, then their beliefs about each player's actions are in equilibrium. It is not that their actions constitute an equilibrium, but that their beliefs do. The question that naturally arises is: are there any epistemic conditions on the players that would induce them to play according to equilibrium? To answer this we turn to Aumann (1987), where it is stated that if players are rational at every state, then their behavior constitutes a correlated equilibrium. Therefore, in order to obtain an equilibrium behavior, a sufficient condition is not that players be rational, or that they know that they are rational at some particular state, but that their rationality be common knowledge. And if it is common knowledge that all players are rational, then their behavior is not necessarily a Nash equilibrium, but a correlated equilibrium.

Theorem 4 *Let G be a strategic game, and let $\langle (\Omega, \pi), (\mathcal{P}_i)_{i \in N}, (\sigma_i)_{i \in N} \rangle$ be a correlated strategy profile for G . If each player is rational at each state of the world, then $\langle (\Omega, \pi), (\mathcal{P}_i)_{i \in N}, (\sigma_i)_{i \in N} \rangle$ is a correlated equilibrium.*

Proof Let $\tau_i : \Omega \rightarrow A_i$ be a function that is measurable with respect to $\mathcal{P}_i$. Since i is Bayes rational at ω

$$E(u_i(\sigma) | \mathcal{P}_i)(\omega) \geq E(u_i(\sigma_{-i}, a_i) | \mathcal{P}_i)(\omega) \quad \forall a_i \in A_i.$$

That is,

$$\begin{aligned} & \sum_{\omega' \in \mathcal{P}_j(\omega)} \frac{\pi(\omega')}{\pi(\mathcal{P}_i(\omega))} u_i(\sigma_{-i}(\omega'), \sigma_i(\omega')) \\ & \geq \sum_{\omega' \in \mathcal{P}_j(\omega)} \frac{\pi(\omega')}{\pi(\mathcal{P}_i(\omega))} u_i(\sigma_{-i}(\omega'), a_i) \quad \forall a_i \in A_i. \end{aligned}$$

In particular, for $a_i = \tau(\omega) = \tau(\omega')$ for all $\omega' \in \mathcal{P}_j(\omega)$,

$$\begin{aligned}
& \sum_{\omega' \in p_j(\omega)} \frac{\pi(\omega')}{\pi(\mathcal{P}_i(\omega))} u_i(\sigma_{-i}(\omega'), \sigma_i(\omega')) \\
& \geq \sum_{\omega' \in p_j(\omega)} \frac{\pi(\omega')}{\pi(\mathcal{P}_i(\omega))} u_i(\sigma_{-i}(\omega'), \tau(\omega')).
\end{aligned}$$

Multiplying both sides by $\pi(\omega)$ and adding over all the elements of $\mathcal{P}_i$ we get

$$\begin{aligned}
& \sum_{\omega \in \Omega} \pi(\omega) u_i(\sigma_{-i}(\omega), \sigma_i(\omega)) \\
& \geq \sum_{\omega \in \Omega} \pi(\omega) u_i(\sigma_{-i}(\omega), \tau(\omega)).
\end{aligned}$$

□

Bayesian Games

Thus far, we have considered static games, which are objects of the form $\langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$. Although these games have many applications, they are not readily suitable for the analysis of situations involving asymmetric information. Indeed, an implicit assumption behind the definition of a static game is that all players have the same information about the relevant aspects of the situation. In particular, all players have the same information about the sets of actions and preferences of all players. A static game seems suitable to model strategic interactions like the prisoner's dilemma, rock scissors paper, and even chess. At the time they choose their actions, all the players have exactly the same information. There might be what is called strategic uncertainty, namely, uncertainty about what the players will do, but there is no uncertainty about the rules of the game and about the preferences of the players. But how would one translate a game of cards like bridge or poker into a static game? In a game of cards, at the time of choosing his actions, each player knows the cards he holds in his hand, but does not know the cards of his opponents. He only has a belief about the cards held by his opponents. In order to make a sound choice, a player will try to predict the actions of his opponents, but for this it is crucial to use his beliefs about the cards they hold. For the same reason, his opponents should use their beliefs

about their own opponents' cards in order to make a sound choice. Thus, the beliefs about the cards held by each player should be part of a description of a game with asymmetric information. Further, in order to predict his opponents' actions, a player also needs to assess his opponents' beliefs about his own cards. This seems to induce an intractable infinite regress of beliefs, and beliefs about beliefs. Harsanyi (1967) provided the basic structure to describe and analyze strategic situations where players are asymmetrically informed. This structure is called a Bayesian game.

Definition 11 A *Bayesian Game* is a system $\langle N, (\Omega, \mu), (A_i, \mathcal{P}_i, u_i)_{i \in N} \rangle$ where

- N is the set of players
- Ω is the set of states of nature
- μ is the players' common prior belief (a probability measure over the set of states)
- A_i is player i 's set of actions
- $\mathcal{P}_i$ is player i 's information partition (a partition of Ω into sets of positive measure). Each element of the partition is referred to as a player's type.
- $u_i : \times_{j \in N} A_j \times \Omega$ is player i 's Bernoulli utility function (a function over pairs (a, ω) where $a \in A$ and $\omega \in \Omega$, the expected value of which represents the player's preferences among lotteries over the set of such pairs).

The interpretation of a Bayesian game is as follows. The basic uncertainty is represented by the probability space (Ω, μ) of all states of nature and the prior probability over them. Each state represents a realization of all the parametric uncertainty of the model. For instance, in a game of cards, each state represents each of the possible card deals. The information of player $i \in N$ is represented by his information partition $\mathcal{P}_i$. While states in the same element of the partition cannot be distinguished by the player, he can distinguish between states that belong to different partition cells. In a game of cards, for instance, each partition cell represents a particular set of

cards dealt to the player. The probability measure μ represents the players' prior belief about the state of nature. This prior belief will be used along with the information obtained by each player to form beliefs about the other players' information. The set of actions of player i is A_i . Note that there is no loss of generality in assuming that this set does not depend on the state nature. One can always add unavailable actions and assign them intolerable disutility. Finally, u_i is the payoff function that associates to each state of nature and action profile a utility level. Note that since the state of the world is unknown to the player at the time of making his choice, a player faces a lottery for any given action profile. The assumption is that the player evaluates this lottery according to the expected value of u_i with respect to that lottery.

Let $\langle N, (\Omega, \mu), (A_i, \mathcal{P}_i, u_i)_{i \in N} \rangle$ be a Bayesian game. A strategy for player $i \in N$ is a function $\sigma_i : \Omega \rightarrow A_i$ that is measurable with respect to $\mathcal{P}_i$. We denote the set of strategies for player i by $\mathcal{B}_i$. That is, $\mathcal{B}_i = \{\sigma_i : \Omega \rightarrow A_i : \sigma_i \text{ is measurable w.r.t. } \mathcal{P}_i\}$. The interpretation of a strategy in a Bayesian game is the usual one. For each state of nature $\omega \in \Omega$, $\sigma_i(\omega)$ is the action chosen by player i at ω . The measurability requirement imposes that player i 's actions depend only on his information. If player i cannot distinguish between two states of nature, then he must choose the same action at both states. Player i evaluates a profile $\sigma : \Omega \rightarrow A$ of strategies according to the expected value of u_i with respect to μ .

In order to define an equilibrium notion for Bayesian games we follow the same idea used for the definition of a mixed strategy equilibrium. Namely, we translate the Bayesian game into a standard game, and then define an equilibrium of the Bayesian game as the Nash equilibrium of the induced game.

Definition 12 A Bayesian equilibrium of a Bayesian game $\langle N, (\Omega, \mu), (A_i, \mathcal{P}_i, u_i)_{i \in N} \rangle$ is a Nash equilibrium of the strategic game: $\langle N, (\mathcal{B}_i)_{i \in N}, (U_i)_{i \in N} \rangle$ where for each profile $\sigma : \Omega \rightarrow A$ of strategies, $U_i(\sigma) = E_\mu[u_i(\sigma(\omega), \omega)]$ is i 's expected utility with respect to μ .

A Bayesian equilibrium of a Bayesian game is a Nash equilibrium of a properly defined static game. As such, conditions for its existence can be derived from Theorem 1. However, in many situations one is interested in particular kinds of equilibria. Specifically, in the analysis of auctions or of the war attrition, one is often interested in efficient outcomes. In a single object auction, efficient outcomes are characterized by the fact that in equilibrium the object is allocated to the buyer who values it most. According to many standard auction rules, the object goes to the highest bidder. Therefore, in such auctions, to guarantee an efficient outcome, one would need a monotone equilibrium, namely, one in which bidders' bids are higher the higher their valuations for the object are. Athey (2001) shows conditions under which a Bayesian equilibrium exists where strategies are non-decreasing. The crucial conditions are that the players' types can be represented by a one-dimensional variable, and that, fixing a nondecreasing strategy for each of a player's opponents, this player's expected payoffs satisfies a single-crossing property. This single-crossing property roughly says that if a high action is preferred to a low action for a given type t , then the same must be true for all types higher than t . McAdams (2003) extended Athey's result to the case where types and actions are multidimensional and partially ordered.

The Asymmetric Information Version of the War of Attrition

We have seen that, when applied to the war of attrition, as modeled by a standard strategic game or by its mixed extension, the notion of Nash equilibrium does not yield a satisfactory prediction. (The war of attrition was analyzed in Maynard Smith (1974). For an analysis of the asymmetric information version of the war attrition, see Krishna and Morgan (1997).) In the former case all the equilibria involve no fight, and in the latter case the equilibrium dictates a more aggressive behavior to the player who

values the contested object less. In what follows, we analyze the war of attrition as a Bayesian game. That is, we assume that the players are ex-ante symmetric but they have private information about their value for the contested object.

A Bayesian game that represents the war of attrition is given by $\langle N, \Omega, (A_i, \mu_i, \mathcal{P}_i, u_i)_{i \in N} \rangle$ where

- $N = \{1, 2\}$
- $\Omega = [0, \infty)^2 = \{(v_1, v_2) : 0 \leq v_i < \infty, i = 1, 2\}$
- $A_i = [0, \infty)$ for $i = 1, 2$
- $\mathcal{P}_i(\hat{v}_1, \hat{v}_2) = \{(v_1, v_2) \in \Omega : v_i = \hat{v}_i\}$ for $i = 1, 2$
- $\mu((v_1, v_2) \leq (\hat{v}_1, \hat{v}_2)) = F(\hat{v}_1) \times F(\hat{v}_2)$
- $u_i((a_1, a_2), (v_1, v_2)) = \begin{cases} -a_i & \text{if } a_i \leq a_j \\ v_i - a_j & \text{if } a_i > a_j. \end{cases}$

Here the set of types player i , for $i = 1, 2$, is represented by the player's willingness to fight, v_i . The players willingness to fight are drawn independently from the same distribution F . A state of the world is, therefore, a realization (v_1, v_2) of the players' types, and at that state, each player is informed only of his type. Finally, the utility of a player is his valuation for the prey, if he obtains it, net of the time spent fighting for it. We are interested in a symmetric equilibrium in which both players use a symmetric, strictly increasing strategy $\beta : [0, \infty) \rightarrow [0, \infty)$, where $\beta(v)$ is the time at which a player with willingness to fight v is dictated by the equilibrium to give up. Such an equilibrium would imply that types who value the prey more, are willing to fight more. Further, the probability of observing a fight in equilibrium would not be 0 (in fact, it would be 1.)

It turns out that a symmetric equilibrium strategy is given by

$$\beta(v) = \int_0^v \frac{xf(x)}{1 - F(x)} dx,$$

where f denotes the derivative of F . To see this, assume that player j behaves according to β and that player i chooses to give up at t . Letting z be the type such $\beta(z) = t$, the expected utility of player i from choosing t is

$$U(v_i, z) = \int_0^z (v_i - \beta(y))f(y)dy - \beta(z)(1 - F(z)).$$

Taking derivatives with respect to z , and using the fact that $\beta'(z) = zf(z)/[1 - F(z)]$ we obtain

$$\begin{aligned} \frac{\partial U}{\partial z}(v_i, z) &= v_i f(z) - \beta'(z)(1 - F(z)) \\ &= (v_i - z)f(z), \end{aligned}$$

which is positive for $z < v_i$, and negative for $z > v_i$. As a result, the expected utility of player i with willingness to pay v_i is maximized at $z = v_i$, which implies that the optimal choice is $\beta(v_i)$.

Thus, modeling the war of attrition as an asymmetric game has allowed us to find an equilibrium in which players with higher willingness to fight fight more, and there is a non-negligible probability of observing a fight.

Evolutionary Stable Strategies

The notion of the Nash equilibrium concept involves players choosing actions that maximize their payoffs given the choices of the other players. The usual interpretation of a Nash equilibrium is as a pattern of behavior that rational players *should* adopt. However, Nash equilibria are sometimes interpreted more descriptively as patterns of behavior that rational players *do* adopt. Certainly, rationality of players is neither a necessary condition nor a sufficient one for players to play a Nash equilibrium. The relationship between rationality and the various solution concepts is not apparent and has been the focus of an extensive literature (see, for example, Aumann 1987; Aumann 1995; Aumann and Brandenburger 1995; Brandenburger and Dekel 1987) Nonetheless, the notion of a Nash equilibrium evokes the idea of players consciously making choices with the deliberate objective of maximizing their payoffs. It is therefore quite remarkable that a concept almost identical to that of Nash equilibrium has emerged from the biology literature. This concept describes a population equilibrium where unconscious organisms are programmed to choose

actions with no deliberate aim. In this equilibrium, members of the population meet at random over and over again to interact. At each interaction, these players act in a pre-programmed way and the result their actions is a gain in biological fitness. Fitness is a concept related to the reproductive value or survival capacity of an organism. In a temporary equilibrium, the fitness gains are such that the proportions of individuals that choose each one of the possible actions remain constant. However, this temporary equilibrium may be disturbed by the appearance of a mutation, which is a new kind of behavior. This mutation may upset the temporary equilibrium if its fitness gains are such that the new behavior spreads over the population. Alternatively, if the fitness gains of the original population outweigh those of the mutation, then the new behavior will fail to propagate and will eventually disappear. In a population equilibrium, the interaction of any mutant with the whole population awards the mutant insufficient fitness gains, and as a result the mutants disappear. The notion of a population equilibrium is formalized by means of the concept of an evolutionary stable strategy, introduced by Maynard Smith and Price (1973).

In what follows we restrict our attention to symmetric two-player games. So let $G = \langle \{1, 2\}, \{A_1, A_2\}, \{u_1, u_2\} \rangle$ be a game such that $A_1 = A_2 = \bar{A}$, and such that for all $a, b \in \bar{A}$, $u_1(a, b) = u_2(b, a)$. An evolutionary stable strategy is an action in $\bar{A}$ such that if all members of the population were to choose that action, no sufficiently small proportion of mutants choosing an alternative action would succeed in invading the population. Alternatively, an evolutionary stable strategy is an action in $\bar{A}$ such that if all the members of the population were to choose that action, the population would reject all sufficiently small mutations involving a different action.

More specifically, suppose that all members of the population are programmed to choose $a \in \bar{A}$, and then a proportion ε of the population mutates and adopts action $b \in \bar{A}$. In that case, the probability that a given member of the population meets a mutant is ε , while the probability of meeting a member that plays a is $1 - \varepsilon$. Therefore, the mutation will not propagate and will vanish if the expected payoff of a mutant is less than the

expected payoff of a member of the majority. Otherwise it will propagate. This leads to the following definition.

Definition 13 An action $a \in \bar{A}$ is an evolutionary stable strategy of G if there is an $\bar{\varepsilon} \in (0, 1)$ such that for all $\varepsilon \in (0, \bar{\varepsilon})$, and for all $b \in \bar{A}$

$$(1 - \varepsilon)u_1(a, a) + \varepsilon u_1(a, b) > (1 - \varepsilon)u_1(b, a) + \varepsilon u_1(b, b). \quad (10)$$

The following result shows that the concept of an evolutionary stable strategy is very close to the notion of a Nash equilibrium.

Proposition 5 If $a \in \bar{A}$ is an evolutionary stable strategy of G , then (a, a) is a Nash equilibrium. And if (a, a) is a strict Nash equilibrium then a is an evolutionary stable strategy.

Proof If $u_1(a, a) > u_1(b, a)$ for all $b \in \bar{A} \setminus \{a\}$, then inequality (10) holds for all sufficiently small $\varepsilon > 0$. If $u_1(b, a) > u_1(a, a)$ for some $b \in \bar{A}$, the reverse inequality holds for all sufficiently small ε . $\square$

Future Directions

Static games have been shown to be a useful framework for analyzing and understanding many situations that involve strategic interaction. At present, a large body of literature is available that develops various solution concepts, some which are refinements of Nash equilibrium and some of which are coarsenings of it. Nonetheless, several areas for future research remain. One is the application of the theory to particular games to better understand the situations they model, for example auctions. In many markets trade is conducted by auctions of one kind or another, including markets for small domestic products as well as some centralized electricity markets where generators and distributors buy and sell electric power on a daily basis. Also, auctions are used to allocate large amounts of valuable spectrum

among telecommunication companies. It would be interesting to calculate the equilibria of many real life auctions. Simultaneously, future research should also focus on the design of auctions whose equilibria have certain desirable properties.

Another future direction would be to empirically and experimentally test the theory. The various equilibrium concepts predict certain kinds of behavior in certain games. Our confidence in the predictive and explanatory power of the theory depends on its performance in the field and in the laboratory. Moreover, the experimental and empirical results should provide valuable feedback for further development of the theory. Although some valuable experimental and empirical tests have already been performed (see McKelvey and Palfrey 1992; O'Neill 1987; Palacios-Huerta 1 2003; Walker and Wooders 2001 to name a few), the empirical aspect of game theory in general, and of static games in particular, remains underdeveloped.

Bibliography

- Athey S (2001) Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica* 69;861–889
- Aumann RJ (1974) Subjectivity and correlation in randomized strategies. *J Math Econ* 1;67–96
- Aumann RJ (1987) Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55;1–18
- Aumann RJ (1995) Backward induction and common knowledge of rationality. *Games Econ Behav* 8;6–19
- Aumann RJ, Brandenburger A (1995) Epistemic conditions for Nash equilibrium. *Econometrica* 63;1161–1180
- Binmore K (2007) *Playing for real*. Oxford University Press, New York
- Brandenburger A, Dekel E (1987) Rationalizability and correlated equilibria. *Econometrica* 55;1391–1402
- Crawford V (1990) Equilibrium without independence. *J Econ Theory* 50;127–154
- Fudenberg D, Tirole J (1991) *Game theory*. MIT Press, Cambridge
- Geanakoplos J (2003) Nash and Walras equilibrium via Brouwer. *Economic Theory* 21;585–603
- Glicksberg IL (1952) A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points. *Proc Am Math Soc* 3;170–174
- Harsanyi J (1967) Games with incomplete information played by 'Bayesian' players. *Manag Sci Part I* 14;159–82; *Part II* 14;320–334; *Part III* 14;486–502
- Krishna V, Morgan J (1997) An analysis of the war of attrition and the all-pay auction. *J Econ Theory* 72;343–362
- McAdams D (2003) Isotone equilibrium in games of incomplete information. *Econometrica* 71;1191–1214
- McKelvey RD, Palfrey TR (1992) An experimental study of the centipede game. *Econometrica* 60;803–836
- Nash JF (1950) Equilibrium points in n -person games. *Proc Natl Acad Sci U S A* 36;48–49
- O'Neill B (1987) Nonmetric test of the minimax theory of two-person zero-sum games. *Proc Natl Acad Sci* 84;2106–2109
- Osborne MJ (2004) *An introduction to game theory*. Oxford University Press, New York
- Osborne MJ, Rubinstein A (1994) *A course in game theory*. MIT Press, Cambridge
- Palacios-Huerta I (2003) Professionals play minimax. *Rev Econ Stud* 70;395–415
- Smith MJ (1974) The theory of games and the evolution of animal conflicts. *J Theor Biol* 47;209–221
- Smith MJ, Price GR (1973) The logic of animal conflict. *Nature* 246: 15–18
- von Neumann J (1928) Zur Theorie der Gesellschaftsspiele. *Math Ann* 100;295–320
- von Neumann J (1959) On the theory of games of strategy. In: Tucker AW, Luce RD (eds) *Contributions to the theory of games*, vol IV. Princeton University Press, Princeton
- Walker M, Wooders J (2001) Minimax play at Wimbledon. *Am Econ Rev* 91;1521–1538

Correlated Equilibria and Communication in Games

Françoise Forges
Ceremade, Université Paris-Dauphine, Paris,
France

Article Outline

Definition of the Subject
Introduction
Correlated Equilibrium: Definition and Basic Properties
Canonical Representation
Duality and Existence
Geometric Properties
Complexity
Foundations
Correlated Equilibrium and Communication
Correlated Equilibrium in Bayesian Games
Strategic Form Correlated Equilibrium
Communication Equilibrium
Correlated Equilibrium, Communication Equilibrium, and Cheap Talk
Related Topics and Future Directions
Bibliography

Glossary

Bayesian game An interactive decision problem consisting of a set of n players, a set of types for every player, a probability distribution which accounts for the players' beliefs over each other's types, a set of actions for every player, and a *von Neumann-Morgenstern utility function* defined over n -tuples of types and actions for every player.

Nash equilibrium In an n -person strategic form game, a strategy n -tuple from which unilateral deviations are not profitable.

Pure strategy (or simply strategy) A mapping which, in an interactive decision problem,

associates an action with the information of a player whenever this player can make a choice.

Sequential equilibrium A refinement of the *Nash equilibrium* for n -person multistage interactive decision problems, which can be loosely defined as a strategy n -tuple together with beliefs over past information for every player, such that every player maximizes his expected utility given his beliefs and the others' strategies, with the additional condition that the beliefs satisfy (possibly sophisticated) Bayes updating given the strategies.

Strategic (or normal) form game An interactive decision problem consisting of a set of n players, a set of strategies for every player, and a (typically, *von Neumann-Morgenstern*) *utility function* defined over n -tuples of strategies for every player.

Utility function A real valued mapping over a set of outcomes which reflects the preferences of an individual by associating a utility level (a "payoff") with every outcome.

von Neumann-Morgenstern utility function A *utility function* which reflects the individual's preferences over lotteries. Such a utility function is defined over outcomes and can be extended to any lottery λ by taking expectation with respect to λ .

Definition of the Subject

The correlated equilibrium is a game theoretic solution concept. It was proposed by Aumann (1974, 1987) in order to capture the strategic correlation opportunities that the players face when they take into account the extraneous environment in which they interact. The notion is illustrated in section "[Introduction](#)." A formal definition is given in section "[Correlated Equilibrium: Definition and Basic Properties](#)." The correlated equilibrium also appears as the appropriate solution concept if preplay communication is allowed between the players. As

shown in section “[Correlated Equilibrium and Communication](#),” this property can be given several precise statements according to the constraints imposed on the players’ communication, which can go from plain conversation to exchange of messages through noisy channels. Originally designed for static games with complete information, the correlated equilibrium applies to any strategic form game. It is geometrically and computationally more tractable than the better known Nash equilibrium. The solution concept has been extended to dynamic games, possibly with incomplete information. As an illustration, we define in detail the communication equilibrium for Bayesian games in section “[Correlated Equilibrium in Bayesian Games](#).”

Introduction

Example

Consider the two-person game known as “chicken,” in which each player i can take a “pacific” action (denoted as p^i) or an “aggressive” action (denoted as a^i):

	p^2	a^2
p^1	(8, 8)	(3, 10)
a^1	(10, 3)	(0, 0)

The interpretation is that player 1 and player 2 simultaneously choose an action and then get a payoff, which is determined by the pair of chosen actions according to the previous matrix. If both players are pacific, they both get 8. If both are aggressive, they both get 0. If one player is aggressive and the other is pacific, the aggressive player gets 10 and the pacific one gets 3. This game has two pure Nash equilibria (p^1, a^2) , (a^1, p^2) and one mixed Nash equilibrium in which both players choose the pacific action with probability $3/5$, resulting in the expected payoff 6 for both players. A possible justification for the latter solution is that the players make their choices as a function of independent extraneous random signals. The assumption of independence is strong. Indeed, there may be no way to prevent the players’ signals from being correlated.

Consider a random signal which has no effect on the players’ payoffs and takes three possible values: low, medium, or high, occurring each with probability $1/3$. Assume that, before the beginning of the game, player 1 distinguishes whether the signal is high or not, while player 2 distinguishes whether the signal is low or not. The relevant interactive decision problem is then the extended game in which the players can base their action on the private information they get on the random signal, while the payoffs only depend on the players’ actions. In this game, suppose that player 1 chooses the aggressive action when the signal is high and the pacific action otherwise. Similarly, suppose that player 2 chooses the aggressive action when the signal is low and the pacific action otherwise. We show that these strategies form an equilibrium in the extended game. Given player 2’s strategy, assume that player 1 observes a high signal. Player 1 deduces that the signal cannot be low so that player 2 chooses the pacific action; hence, player 1’s best response is to play aggressively. Assume now that player 1 is informed that the signal is not high; he deduces that, with probability $1/2$, the signal is medium (i.e., not low) so that player 2 plays pacific and, with probability $1/2$, the signal is low so that player 2 plays aggressive. The expected payoff of player 1 is 5.5 if he plays pacific and 5 if he plays aggressive; hence, the pacific action is a best response. The equilibrium conditions for player 2 are symmetric. To sum up, the strategies based on the players’ private information form a Nash equilibrium in the extended game in which an extraneous signal is first selected. We shall say that these strategies form a “correlated equilibrium.” The corresponding probability distribution over the players’ actions is

$$\begin{array}{cc}
 & \begin{matrix} p^2 & a^2 \end{matrix} \\
 \begin{matrix} p^1 \\ a^1 \end{matrix} & \begin{pmatrix} \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & 0 \end{pmatrix}
 \end{array} \tag{1}$$

and the expected payoff of every player is 7. This probability distribution can be used directly to make private recommendations to the players

before the beginning of the game (see the section “[Canonical Representation](#)” below).

Correlated Equilibrium: Definition and Basic Properties

Definition

A game in strategic form $G = (N, (\Sigma^i)_{i \in N}, (u^i)_{i \in N})$ consists of a set of players N together with, for every player $i \in N$, a set of strategies (for instance, a set of actions) Σ^i and a (von Neumann-Morgenstern) utility function $u^i : \Sigma \rightarrow \mathbb{R}$, where $\Sigma = \prod_{j \in N} \Sigma^j$ is the set of all strategy profiles. We assume that the sets N and Σ^i , $i \in N$, are finite.

A correlation device $d = (\Omega, q, (\mathcal{P}^i)_{i \in N})$ is described by a finite set of signals Ω , a probability distribution q over Ω , and a partition $\mathcal{P}^i$ of Ω for every player $i \in N$. Since Ω is finite, the probability distribution q is just a real vector $q = (q(\omega))_{\omega \in \Omega}$ such that $q(\omega) \geq 0$ and $\sum_{\omega \in \Omega} q(\omega) = 1$.

From G and d , we define the extended game G_d as follows:

- ω is chosen in Ω according to q .
- Every player i is informed of the element $P^i(\omega)$ of $\mathcal{P}^i$ which contains ω .
- G is played: every player i chooses a strategy σ^i in Σ^i and gets the utility $u^i(\sigma)$, $\sigma = (\sigma^j)_{j \in N}$.

A (pure) strategy for player i in G_d is a mapping $\alpha^i : \Omega \rightarrow \Sigma^i$ which is $\mathcal{P}^i$ -measurable, such that $\alpha^i(\omega') = \alpha^i(\omega)$ if $\omega' \in P^i(\omega)$. The interpretation is that, in G_d , every player i chooses his strategy σ^i as a function of his private information on the random signal ω which is selected before the beginning of G .

According to Aumann (1974), a *correlated equilibrium* of G is a pair (d, α) , which consists of a correlation device $d = (\Omega, q, (\mathcal{P}^i)_{i \in N})$ and a Nash equilibrium $\alpha = (\alpha^i)_{i \in N}$ of G_d . The equilibrium conditions of every player i , conditionally on his private information, can be written as

$$\begin{aligned} & \sum_{\omega' \in P^i(\omega)} q(\omega' | P^i(\omega)) u^i(\alpha(\omega')) \\ & \geq \sum_{\omega' \in P^i(\omega)} q(\omega' | P^i(\omega)) u^i(\tau^i, \alpha^{-i}(\omega')), \quad (2) \\ & \forall i \in N, \forall \tau^i \in \Sigma^i, \forall \omega \in \Omega : q(\omega) > 0, \end{aligned}$$

where $\alpha^{-i} = (\alpha^j)_{j \neq i}$.

A mixed Nash equilibrium $\rho = (\rho^i)_{i \in N}$ of G can be viewed as a correlated equilibrium of G . By definition, every ρ^i is a probability distribution over Σ^i , the finite set of pure strategies of player i . Let us consider the correlation device $d = (\Omega, q, (\mathcal{P}^i)_{i \in N})$ in which $\Omega = \Sigma = \prod_{j \in N} \Sigma^j$, q is the product probability distribution induced by the mixed strategies (i.e., $q((\sigma^j)_{j \in N}) = \prod_{j \in N} \rho^j(\sigma^j)$, and, for each i , $\mathcal{P}^i$ is the partition of Ω generated by Σ^i (i.e., for $\omega, v \in \Omega$, $v \in P^i(\omega) \Leftrightarrow v^i = \omega^i$). Let $\alpha^i : \Sigma \rightarrow \Sigma^i$ be the projection over Σ^i (i.e., $\alpha^i(\sigma) = \sigma^i$). The correlation device d and the strategies α^i defined in this way form a correlated equilibrium. As we shall see below, this correlated equilibrium is “canonical.”

Canonical Representation

A *canonical* correlated equilibrium of G is a correlated equilibrium in which $\Omega = \Sigma = \prod_{j \in N} \Sigma^j$, while for every player i , the partition $\mathcal{P}^i$ of Σ is generated by Σ^i and $\alpha^i : \Sigma \rightarrow \Sigma^i$ is the projection over Σ^i . A canonical correlated equilibrium is thus fully specified by a probability distribution q over Σ . A natural interpretation is that a mediator selects $\sigma = (\sigma^j)_{j \in N}$ according to q and privately recommends σ^i to player i , for every $i \in N$. The players are not forced to obey the mediator, but σ is selected in such a way that player i cannot benefit from deviating unilaterally from the recommendation σ^i , i.e., $\tau^i = \sigma^i$ maximizes the conditional expectation of player i 's payoff $u^i(\tau^i, \sigma^{-i})$ given the recommendation σ^i . A probability distribution q over Σ thus defines a canonical correlated equilibrium if and only if it satisfies the following linear inequalities:

$$\begin{aligned}
& \sum_{\sigma^{-i} \in \Sigma^{-i}} \mathbf{q}(\sigma^{-i} | \sigma^i) u^i(\sigma^i, \sigma^{-i}) \\
& \geq \sum_{\sigma^{-i} \in \Sigma^{-i}} \mathbf{q}(\sigma^{-i} | \sigma^i) u^i(\tau^i, \sigma^{-i}), \\
& \forall i \in N, \forall \sigma^i \in \Sigma^i : \mathbf{q}(\sigma^{-i}) > 0, \forall \tau^i \in \Sigma^i
\end{aligned}$$

or, equivalently,

$$\begin{aligned}
& \sum_{\sigma^{-i} \in \Sigma^{-i}} \mathbf{q}(\sigma^i, \sigma^{-i}) u^i(\sigma^i, \sigma^{-i}) \\
& \geq \sum_{\sigma^{-i} \in \Sigma^{-i}} \mathbf{q}(\sigma^i | \sigma^{-i}) u^i(\tau^i, \sigma^{-i}), \quad (3) \\
& \forall i \in N, \forall \sigma^i, \tau^i \in \Sigma^i
\end{aligned}$$

The equilibrium conditions can also be formulated ex ante:

$$\begin{aligned}
& \sum_{\sigma \in \Sigma} \mathbf{q}(\sigma) u^i(\sigma) \geq \sum_{\sigma \in \Sigma} \mathbf{q}(\sigma) u^i(\alpha^i(\sigma^i), \sigma^{-i}), \\
& \forall i \in N, \forall \alpha^i : \Sigma^i \rightarrow \Sigma^i
\end{aligned}$$

The following result is an analog of the “revelation principle” in mechanism design (see, e.g., Myerson 1982): let $(\alpha, \mathbf{d})$ be a correlated equilibrium associated with an arbitrary correlation device $d = (\Omega, \mathbf{q}, (\mathcal{P}^i)_{i \in N})$. The corresponding “correlated equilibrium distribution,” namely, the probability distribution induced over Σ by $\mathbf{q}$ and α , defines a canonical correlated equilibrium. For instance, in the introduction, Eq. 1 describes a canonical correlated equilibrium.

Duality and Existence

From the linearity of Eq. 3, duality theory can be used to study the properties of correlated equilibria, in particular to prove their existence without relying on Nash’s (1951) theorem and its fixed point argument (recall that every mixed Nash equilibrium is a correlated equilibrium). Hart and Schmeidler (1989) establish the existence of a correlated equilibrium by constructing an auxiliary two-person zero-sum game and applying the minimax theorem. Nau and McCardle (1990) derive another elementary proof of existence from an extension of the “no

arbitrage opportunities” axiom that underlies subjective probability theory. They introduce jointly coherent strategy profiles, which do not expose the players as a group to arbitrage from an outside observer. They show that a strategy profile is jointly coherent if and only if it occurs with positive probability in some correlated equilibrium. From a technical point of view, both proofs turn out to be similar. Myerson (1997) makes further use of the linear structure of correlated equilibria by introducing dual reduction, a technique to replace a finite game with a game with fewer strategies, in such a way that any correlated equilibrium of the reduced game induces a correlated equilibrium of the original game.

Geometric Properties

As Eq. 3 is a system of linear inequalities, the set of all correlated equilibrium distributions is a convex polytope. Nau et al. (2004) show that if it has “full” dimension (namely, dimension $|\Sigma| - 1$), then all Nash equilibria lie on its relative boundary. Viossat (2006) characterizes in addition the class of games whose correlated equilibrium polytope contains a Nash equilibrium in its relative interior. Interestingly, this class of games includes two-person zero-sum games but is not defined by “strict competition” properties. In two-person games, all extreme Nash equilibria are also extreme correlated equilibria (Evangelista and Raghavan 1996; Gomez-Canovas et al. 1999); this result does not hold with more than two players. Finally, Viossat (2008) proves that having a unique correlated equilibrium is a robust property, in the sense that the set of n -person games with a unique correlated equilibrium is open. The same is not true for Nash equilibrium (unless $n = 2$).

Complexity

From Eq. 3, correlated equilibria can be computed by linear programming methods. Gilboa

and Zemel (1989) show more precisely that the complexity of standard computational problems is “NP-hard” for the Nash equilibrium and *polynomial* for the correlated equilibrium. Examples of such problems are “Does the game G have a Nash (resp., correlated) equilibrium which yields a payoff greater than r to every player (for some given number r)?” and “Does the game G have a unique Nash (resp., correlated) equilibrium?” Papadimitriou (2005) develops a polynomial-time algorithm for finding correlated equilibria, which is based on a variant of the existence proof of Hart and Schmeidler (1989).

Foundations

By reinterpreting the previous canonical representation, Aumann (1987) proposes a decision theoretic foundation for the correlated equilibrium in games with complete information, in which Σ^i for $i \in N$ stands merely for a set of *actions* of player i . Let Ω be the space of all states of the world; an element ω of Ω thus specifies all the parameters which may be relevant to the players’ choices. In particular, the action profile in the underlying game G is part of the state of the world. A partition $\mathcal{P}^i$ describes player i ’s information on Ω . In addition, every player i has a prior belief, i.e., a probability distribution q^i over Ω . Formally, the framework is similar as above except that the players possibly hold different beliefs over Ω . Let $\alpha^i(\omega)$ denote player i ’s action at ω ; a natural assumption is that player i knows the action he chooses, namely, that α^i is $\mathcal{P}^i$ -measurable. According to Aumann (1987), player i is *Bayes rational* at ω if his action $\alpha^i(\omega)$ maximizes his expected payoff (with respect to q^i) given his information $P^i(\omega)$. Note that this is a separate rationality condition for every player, not an equilibrium condition. Aumann (1987) proves the following result: *under the common prior assumption (namely, $q^i = q$, $i \in N$), if every player is Bayes rational at every state of the world, the distribution of the corresponding action profile α is a correlated equilibrium distribution*. The key to this decision

theoretic foundation of the correlated equilibrium is that, under the common prior assumption, Bayesian rationality amounts to Eq. 2.

If the common prior assumption is relaxed, the previous result still holds, with *subjective* prior probability distributions, for the *subjective* correlated equilibrium which was also introduced by Aumann (1974). The latter solution concept is defined in the same way as above, by considering a device $(\Omega, (q^i)_{i \in N}, (P^i)_{i \in N})$, with a probability distribution q^i for every player i , and by writing Eq. 2 in terms of q^i instead of q . Brandenburger and Dekel (1987) show that (a refinement of) the subjective correlated equilibrium is equivalent to (correlated) rationalizability, another well-established solution concept which captures players’ minimal rationality. Rationalizable strategies reflect that the players commonly know that each of them makes an optimal choice given some belief. Nau and McCardle (1991) reconcile objective and subjective correlated equilibrium by proposing the no arbitrage principle as a unified approach to individual and interactive decision problems. They argue that the objective correlated equilibrium concept applies to a game that is revealed by the players’ choices, while the subjective correlated equilibrium concept applies to the “true game”; both lead to the same set of jointly coherent outcomes.

Correlated Equilibrium and Communication

As seen in the previous section, correlated equilibria can be achieved in practice with the help of a mediator and emerge in a Bayesian framework embedding the game in a full description of the world. Both approaches require to extend the game by taking into account information which is not generated by the players themselves. Can the players reach a correlated equilibrium without relying on any extraneous correlation device, by just communicating with each other before the beginning of the game?

Consider the game of “chicken” presented in the introduction. The probability distribution

$$\begin{array}{cc}
 & \begin{array}{c} p^2 \\ a^2 \end{array} \\
 \begin{array}{c} p^1 \\ a^1 \end{array} & \begin{array}{cc} 0 & \frac{1}{2} \\ \frac{1}{2} & 0 \end{array}
 \end{array} \quad (4)$$

describes a correlated equilibrium, which amounts to choosing one of the two pure Nash equilibria, with equal probability. Both players get an expected payoff of 6.5. Can they safely achieve this probability distribution if no mediator tosses a fair coin for them? The answer is positive, as shown by Aumann et al. (1968). Assume that before playing “chicken,” the players independently toss a coin and simultaneously reveal to each other whether heads or tails obtains. Player 1 tells player 2 “ h^1 ” or “ t^1 ” and, at the same time, player 2 tells player 1 “ h^2 ” or “ t^2 .” If both players use a fair coin, reveal correctly the result of the toss and play (p^1, a^2) if both coins fell on the same side (i.e., if (h^1, h^2) or (t^1, t^2) is announced) and (a^1, p^2) otherwise (i.e., if (h^1, t^2) or (t^1, h^2) is announced); they get the same effect as a mediator using Eq. 4. Furthermore, none of them can gain by unilaterally deviating from the described strategies, even at the randomizing stage: the two relevant outcomes, $[(h^1, h^2)$ or $(t^1, t^2)]$ and $[(h^1, t^2)$ or $(t^1, h^2)]$, happen with probability $1/2$ provided that one of the players reveals the toss of a fair coin. This procedure is known as a “jointly controlled lottery.” An important feature of the previous example is that, in the correlated equilibrium described by Eq. 4, the players know each other’s recommendation. Hence, they can easily reproduce (Eq. 4) by exchanging messages that they have selected independently. In the correlated equilibrium described by the probability distribution (1), the private character of recommendations is crucial to guarantee that (p^1, p^2) be played with positive probability. Hence, one cannot hope that a simple procedure of direct preplay communication be sufficient to generate (Eq. 1). However, the fact that direct communication is necessarily public is typical of two-person games.

Given the game $G = (N, (\Sigma^i)_{i \in N}, (u^i)_{i \in N})$, let us define a (bounded) “cheap talk” extension $\text{ext}(G)$ of G as a game in which T stages of costless, unmediated preplay communication are allowed before G is played. More precisely, let M_t^i

be a finite set of messages for player i , $i \in N$, at stage t , $t = 1, 2, \dots, T$; at every stage t of $\text{ext}(G)$, every player i selects a message in M_t^i ; these choices are made simultaneously before being revealed to a subset of players at the end of stage t . The rules of $\text{ext}(G)$ thus determine a set of “senders” for every stage t (those players i for whom M_t^i contains more than one message) and a set of “receivers” for every stage t . The players perfectly recall their past messages. After the communication phase, they choose their strategies (e.g., their actions) as in G ; they are also rewarded as in G , independently of the preplay phase, which is thus “cheap.” Communication has an indirect effect on the final outcome in G , since the players make their decisions as a function of the messages that they have exchanged. Specific additional assumptions are often made on $\text{ext}(G)$, as we will see below.

Let us fix a cheap talk extension $\text{ext}(G)$ of G and a Nash equilibrium of $\text{ext}(G)$. As a consequence of the previous definitions, the distribution induced by this Nash equilibrium over Σ defines a correlated equilibrium of G (this can be proved in the same way as the canonical representation of correlated equilibria stated in section “Correlated Equilibrium: Definition and Basic Properties”). The question raised in this section is whether the reverse holds.

If the number of players is two, the Nash equilibrium distributions of cheap talk extensions of G form a subset of the correlated equilibrium distributions: the convex hull of Nash equilibrium distributions. Indeed, the players have both the same information after any direct exchange of messages. Conversely, by performing repeated jointly controlled lotteries like in the example above, the players can achieve any convex combination (with rational weights) of Nash equilibria of G as a Nash equilibrium of a cheap talk extension of G . The restriction on probability distributions whose components are rational numbers is only needed as far as we focus on *bounded* cheap talk extensions.

Bárány (1992) establishes that, if the number of players of G is at least four, every (rational) correlated equilibrium distribution of G can be realized as a Nash equilibrium of a cheap talk extension $\text{ext}(G)$, provided that $\text{ext}(G)$ allows the

players to publicly check the record of communication under some circumstances. The equilibria of $\text{ext}(G)$ constructed by Bárány involve that a receiver gets the same message from two different senders; the message is nevertheless not public, thanks to the assumption on the number of players. At every stage of $\text{ext}(G)$, every player can ask for the revelation of all past messages, which are assumed to be recorded. Typically, a receiver can claim that the two senders' messages differ. In this case, the record of communication surely reveals that either one of the senders or the receiver himself has cheated; the deviator can be punished (at his minimax level in G) by the other players.

The punishments in Bárány's (1992) Nash equilibria of $\text{ext}(G)$ need not be credible threats. Instead of using double senders in the communication protocols, Ben-Porath (1998, 2003) proposes a procedure of random monitoring, which prescribes a given behavior to every player in such a way that unilateral deviations can be detected with probability arbitrarily close to 1. This procedure applies if there are at least three players, which yields an analog of Bárány's result already in this case. If the number of players is exactly three, Ben-Porath (2003) needs to assume, as Bárány (1992), that public verification of the record of communication is possible in $\text{ext}(G)$ (see Ben-Porath 2006). However, Ben-Porath concentrates on (rational) correlated equilibrium distributions which allow for strict punishment on a Nash equilibrium of G ; he constructs *sequential* equilibria which generate these distributions in $\text{ext}(G)$, thus dispensing with incredible threats. At the price of raising the number of players to five or more, Gerardi (2004) proves that every (rational) correlated equilibrium distribution of G can be realized as a sequential equilibrium of a cheap talk extension of G which does not require any message recording. For this, he builds protocols of communication in which the players base their decisions on majority rule, so that no punishment is necessary.

We have concentrated on two extreme forms of communication: mediated communication, in which a mediator performs lotteries and sends private messages to the players, and cheap talk, in which the players just exchange messages. Many intermediate schemes of communication are obviously conceivable. For instance, Lehrer (1996)

introduces (possibly multistage) "mediated talk": the players send private messages to a mediator, but the latter can only make deterministic public announcements. Mediated talk captures real-life communication procedures, like elections, especially if it lasts only for a few stages. Lehrer and Sorin (1997) establish that whatever the number of players of G , every (rational) correlated equilibrium distribution of G can be realized as a Nash equilibrium of a single-stage mediated talk extension of G . Ben-Porath (1998) proposes a variant of cheap talk in which the players do not only exchange verbal messages but also "hard" devices such as urns containing balls. This extension is particularly useful in two-person games to circumvent the equivalence between the equilibria achieved by cheap talk and the convex hull of Nash equilibria. More precisely, the result of Ben-Porath (1998) stated above holds for two-person games if the players first check together the content of different urns and then each player draws a ball from an urn that was chosen by the other player, so as to guarantee that one player only knows the outcome of a lottery, while the other one only knows the probabilities of this lottery.

The various extensions of the basic game G considered up to now, with or without a mediator, implicitly assume that the players are fully rational. In particular, they have unlimited computational abilities. By relaxing that assumption, Urbano and Vila (2002) and Dodis et al. (2000) build on earlier results from cryptography so as to implement any (rational) correlated equilibrium distribution through unmediated communication, including in two-person games.

As the previous paragraphs illustrate, the players can modify their initial distribution of information by means of many different communication protocols. Gossner (1998) proposes a general criterion to classify them: a protocol is "secure" if under all circumstances, the players cannot mislead each other nor spy on each other. For instance, given a cheap talk extension $\text{ext}(G)$, a protocol P describes, for every player, a strategy in $\text{ext}(G)$ and a way to interpret his information after the communication phase of $\text{ext}(G)$. P induces a correlation device $d(P)$ (in the sense of section "Correlated Equilibrium: Definition and Basic Properties"). P is secure if, for every game G and every Nash equilibrium α of $G_{d(P)}$, the following procedure is a Nash

equilibrium of $\text{ext}(G)$: communicate according to the strategies described by P in order to generate $d(P)$ and make the final choice, in G , according to α . Gossner (1998) gives a tractable characterization of secure protocols.

Correlated Equilibrium in Bayesian Games

A Bayesian game $\Gamma = (N, (T^i)_{i \in N}, p, (A^i)_{i \in N}, (u^i)_{i \in N})$ consists of a set of players N : for every player $i \in N$, a set of types T^i ; a probability distribution p^i over $T = \prod_{j \in N} T^j$; a set of actions A^i ; and a (von Neumann-Morgenstern) utility function $u^i : T \times A \rightarrow \mathbb{R}$, where $A = \prod_{j \in N} A^j$. For simplicity, we make the common prior assumption: $p^i = p$ for every $i \in N$. All sets are assumed finite. The interpretation is that a virtual move of nature chooses $t = (t^j)_{j \in N}$ according to p ; player i is only informed of his own type t^i ; the players then choose simultaneously an action. We will focus on two possible extensions of Aumann's (1974) solution concept to Bayesian games: the strategic form correlated equilibrium and the communication equilibrium. Without loss of generality, the definitions below are given in "canonical form" (see section "Correlated Equilibrium: Definition and Basic Properties").

Strategic Form Correlated Equilibrium

A (pure) strategy of player i in Γ is a mapping $\sigma^i : T^i \rightarrow A^i$, $i \in N$. The strategic form of Γ is a game $G(\Gamma)$, like the game G considered in section "Correlated Equilibrium: Definition and Basic Properties," with sets of pure strategies $\Sigma_i = A_i^{T_i}$ and utility functions u^i over $\Sigma = \prod_{j \in N} \Sigma^j$ computed as expectations with respect to $p : u^i(\sigma) = E[u^i(t, \sigma(t))]$, with $\sigma(t) = (\sigma^i(t^i))_{i \in N}$. A strategic form correlated equilibrium, or simply a correlated equilibrium, of a Bayesian game Γ is a correlated equilibrium, in the sense of section "Correlated Equilibrium: Definition and Basic Properties," of $G(\Gamma)$. A canonical correlated equilibrium of Γ is thus described by a probability

distribution Q over Σ , which selects an N -tuple of pure strategies $(\sigma^i)_{i \in N}$. This lottery can be thought of as being performed by a mediator who privately recommends σ^i to player i , $i \in N$, before the beginning of Γ , i.e., before (or in any case, independently of) the chance move choosing the N -tuple of types. The equilibrium conditions express that, once he knows his type t^i , player i cannot gain in unilaterally deviating from $\sigma^i(t^i)$.

Communication Equilibrium

Myerson (1982) transforms the Bayesian game Γ into a mechanism design problem by allowing the mediator to collect information from the players before making them recommendations. Following Forges (1986) and Myerson (1986a), a canonical communication device for Γ consists of a system q of probability distributions $q = (q(\cdot|t))_{t \in T}$ over A . The interpretation is that a mediator invites every player i , $i \in N$, to report his type t^i , then selects an N -tuple of actions a according to $q(\cdot|t)$, and privately recommends a^i to player i . The system q defines a communication equilibrium if none of the players can gain by unilaterally lying on his type or by deviating from the recommended action, namely, if

$$\begin{aligned} & \sum_{t^{-i} \in T^{-i}} p(t^{-i}|t^i) \sum_{a \in A} q(a|t) u^i(t, a) \\ & \geq \sum_{t^{-i} \in T^{-i}} p(t^{-i}|t^i) \sum_{a \in A} q(a|s^i, t^{-i}) u^i(t, \alpha^i(a^i), a^{-i}), \\ & \quad \forall i \in N, \forall t^i, s^i \in T^i, \forall \alpha^i : A^i \rightarrow A^i \end{aligned}$$

Correlated Equilibrium, Communication Equilibrium, and Cheap Talk

Every correlated equilibrium of the Bayesian game Γ induces a communication equilibrium of Γ , but the converse is not true, as the following example shows.

Consider the two-person Bayesian game in which $T^1 = \{s^1, t^1\}$, $T^2 = \{t^2\}$, $A^1 = \{a^1, b^1\}$, $A^2 = \{a^2, b^2\}$, $p(s^1) = p(t^1) = \frac{1}{2}$, and payoffs are described by

		a^2	b^2
s^1	a^1	(1, 1)	(-1, -1)
	b^1	(0, 0)	(0, 0)
		a^2	b^2
t^1	a^1	(0, 0)	(0, 0)
	b^1	(-1, -1)	(1, 1)

In this game, the communication equilibrium $q(a^1, a^2 | s^1) = q(b^1, b^2 | t^1) = 1$ yields the expected payoff of 1 to both players. However, the maximal expected payoff of every player in a correlated equilibrium is $1/2$. In order to see this, one can derive the strategic form of the game (in which player 1 has four strategies and player 2 has two strategies). Let us turn to the game in which player 1 can cheaply talk to player 2 just after having learned his type. In this new game, the following strategies form a Nash equilibrium: player 1 truthfully reveals his type to player 2 and plays a^1 if s^1 , b^1 if t^1 ; player 2 chooses a^2 if s^1 , b^2 if t^1 . These strategies achieve the same expected payoffs as the communication equilibrium.

As in section “[Correlated Equilibrium and Communication](#),” one can define cheap talk extensions $\text{ext}(\Gamma)$ of Γ . A wide definition of $\text{ext}(\Gamma)$ involves an ex ante preplay phase, before the players learn their types, and an interim preplay phase, after the players learn their types but before they choose their actions. Every Nash equilibrium of $\text{ext}(\Gamma)$ induces a communication equilibrium of Γ . In order to investigate the converse, namely, whether cheap talk can simulate mediated communication in a Bayesian game, two approaches have been developed. The first one (Forges 1990; Gerardi 2000, 2004; Vida 2007) proceeds in two steps, by reducing communication equilibria to correlated equilibria before applying the results obtained for strategic form games (see section “[Correlated Equilibrium and Communication](#)”). The second approach (Ben-Porath 2003; Krishna 2007) directly addresses the question in a Bayesian game.

By developing a construction introduced for particular two-person games, Forges (1985, 1990) shows that every communication equilibrium outcome of a Bayesian game Γ with at least

four players can be achieved as a correlated equilibrium outcome of a two-stage interim cheap talk extension $\text{ext}_{\text{int}}(\Gamma)$ of Γ . No punishment is necessary in $\text{ext}_{\text{int}}(\Gamma)$: at the second stage, every player gets a message from three senders and uses majority rule if the messages are not identical. Thanks to the underlying correlation device, each receiver is able to privately decode his message. Vida (2007) extends Forges (1990) to Bayesian games with three or even two players. In the proof, he constructs a correlated equilibrium of a long, but almost surely finite, interim cheap talk extension of Γ , whose length depends both on the signals selected by the correlation device and the messages exchanged by the players. No recording of messages is necessary to detect and punish a cheating player. If there are at least four players in Γ , once a communication equilibrium of Γ has been converted into a correlated equilibrium of $\text{ext}_{\text{int}}(\Gamma)$, one can apply Bárány’s (1992) result to $\text{ext}_{\text{int}}(\Gamma)$ in order to transform the correlated equilibrium into a Nash equilibrium of a further, ex ante, cheap talk preplay extension of Γ . Gerardi (2000) modifies this ex ante preplay phase so as to postpone it at the interim stage. This result is especially useful if the initial move of nature in Γ is just a modeling convenience. Gerardi (2004) also extends his result for at least five-person games with complete information (see section “[Correlated Equilibrium and Communication](#)”) to any Bayesian game with full support (i.e., in which all type profiles have positive probability: $p(t) > 0$ for every $t \in T$) by proving that every (rational) communication equilibrium of Γ can be achieved as a *sequential* equilibrium of a cheap talk extension of Γ .

Ben-Porath (2003) establishes that if Γ is a three (or more)-person game with full support, every (rational) communication equilibrium of Γ which strictly dominates a Nash equilibrium of Γ for every type t^i of every player i , $i \in N$, can be implemented as a Nash equilibrium of an interim cheap talk extension of Γ in which public verification of past record is possible (see also Ben-Porath 2006). Krishna (2007) extends Ben-Porath’s (1998) result on two-person games (see section “[Correlated Equilibrium and Communication](#)”) to the incomplete information

framework. The other results mentioned at the end of section “[Correlated Equilibrium and Communication](#)” have also been generalized to Bayesian games (see Gossner 1998; Lehrer and Sorin 1997; Urbano and Vila 2004a).

Related Topics and Future Directions

In this brief entry, we concentrated on two solution concepts: the strategic form correlated equilibrium, which is applicable to any game, and the communication equilibrium, which we defined for Bayesian games. Other extensions of Aumann’s (1974) solution concept have been proposed for Bayesian games, as the agent normal form correlated equilibrium and the (possibly belief invariant) Bayesian solution (see Forges (1993, 2006) for definitions and references). The Bayesian solution is intended to capture the players’ rationality in games with incomplete information in the spirit of Aumann (1987) (see Nau 1992; Forges 1993). Lehrer et al. (2006) open a new perspective in the understanding of the Bayesian solution and other equilibrium concepts for Bayesian games by characterizing the classes of equivalent information structures with respect to each of them. Comparison of information structures, which goes back to Blackwell (1951, 1953) for individual decision problems, was introduced by Gossner (2000) in the context of games, both with complete and incomplete information. In the latter model, information structures basically describe how extraneous signals are selected as a function of the players’ types; two information structures are equivalent with respect to an equilibrium concept if, in every game, they generate the same equilibrium distributions over outcomes.

Correlated equilibria, communication equilibria, and related solution concepts have been studied in many other classes of games, like multistage games (see, e.g., Forges 1986; Myerson 1986a), repeated games with incomplete information (see, e.g., Forges 1985, 1988), and stochastic games (see, e.g., Solan 2001; Solan and Vieille 2002). The study of correlated equilibrium in repeated games with imperfect monitoring, initiated by

Lehrer (1991, 1992), proved to be particularly useful and is still undergoing. Lehrer (1991) showed that if players either are fully informed of past actions or get no information (“standard-trivial” information structure), correlated equilibria are equivalent to Nash equilibria. In other words, all correlations can be generated internally, namely, by the past histories, on which players have differential information. The schemes of internal correlation introduced to establish this result are widely applicable and inspired those of Lehrer (1996) (see section “[Correlated Equilibrium and Communication](#)”). In general repeated games with imperfect monitoring, Renault and Tomala (2004) characterize communication equilibria, but the amount of correlation that the players can achieve in a Nash equilibrium is still an open problem (see, e.g., Gossner and Tomala 2007; Urbano and Vila 2004b for recent advances).

Throughout this entry, we defined a correlated equilibrium as a *Nash* equilibrium of an extension of the game under consideration. The solution concept can be strengthened by imposing some refinement, i.e., further rationality conditions, to the Nash equilibrium in this definition (see, e.g., Dhillon and Mertens 1996; Myerson 1986b). Refinements of communication equilibria have also been proposed (see, e.g., Gerardi 2004; Gerardi and Myerson 2007; Myerson 1986a). Some authors (see, e.g., Milgrom and Roberts 1996; Moreno and Wooders 1996; Ray 1996) have also developed notions of *coalition proof* correlated equilibria, which resist not only to unilateral deviations, as in this entry, but even to multilateral ones. A recurrent difficulty is that, for many of these stronger solution concepts, a useful canonical representation (as derived in section “[Correlated Equilibrium: Definition and Basic Properties](#)”) is not available.

Except for two or three references, we deliberately concentrated on the results published in the game theory and mathematical economics literature, while substantial achievements in computer science would fit in this survey. Both streams of research pursue similar goals but rely on different formalisms and techniques. For instance, computer scientists often make use of cryptographic tools which are not familiar in

game theory. Halpern (2007) gives an idea of recent developments at the interface of computer science and game theory (see in particular the section “Implementing Mediators”) and contains a number of references.

Finally, the assumption of full rationality of the players can also be relaxed. Evolutionary game theory has developed models of learning in order to study the long-term behavior of players with bounded rationality. Many possible dynamics are conceivable to represent more or less myopic attitudes with respect to optimization. Under appropriate learning procedures, which express, for instance, that agents want to minimize the regret of their strategic choices, the empirical distribution of actions converges to correlated equilibrium distributions (see, e.g., Foster and Vohra 1997; Hart and Mas-Colell 2000; Hart 2005 for a survey). However, standard procedures, as the “replicator dynamics,” may even eliminate all the strategies which have positive probability in a correlated equilibrium (see Viossat 2007).

Bibliography

Primary Literature

- Aumann RJ (1974) Subjectivity and correlation in randomized strategies. *J Math Econ* 1:67–96
- Aumann RJ (1987) Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55:1–18
- Aumann RJ, Maschler M, Stearns R (1968) Repeated games with incomplete information: an approach to the nonzero sum case. Reports to the US Arms Control and Disarmament Agency, ST-143, Chapter IV, 117–216 (reprinted In: Aumann RJ, Maschler M (1995) *Repeated Games of Incomplete Information*. M.I.T. Press, Cambridge)
- Bárány I (1992) Fair distribution protocols or how players replace fortune. *Math Oper Res* 17:327–340
- Ben-Porath E (1998) Correlation without mediation: expanding the set of equilibrium outcomes by cheap pre-play procedures. *J Econ Theory* 80:108–122
- Ben-Porath E (2003) Cheap talk in games with incomplete information. *J Econ Theory* 108:45–71
- Ben-Porath E (2006) A correction to “Cheap talk in games with incomplete information”. Mimeo, Hebrew University of Jerusalem, Jerusalem
- Blackwell D (1951) Comparison of experiments. In: *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, University of California Press, Berkeley, pp 93–102
- Blackwell D (1953) Equivalent comparison of experiments. *Ann Math Stat* 24:265–272
- Brandenburger A, Dekel E (1987) Rationalizability and correlated equilibria. *Econometrica* 55:1391–1402
- Dhillon A, Mertens JF (1996) Perfect correlated equilibria. *J Econ Theory* 68:279–302
- Dodis Y, Halevi S, Rabin T (2000) A cryptographic solution to a game theoretic problem. In: *CRYPTO 2000: 20th international cryptology conference*. Springer, Berlin, pp 112–130
- Evangelista F, Raghavan TES (1996) A note on correlated equilibrium. *Int J Game Theory* 25:35–41
- Forges F (1985) Correlated equilibria in a class of repeated games with incomplete information. *Int J Game Theory* 14:129–150
- Forges F (1986) An approach to communication equilibrium. *Econometrica* 54:1375–1385
- Forges F (1988) Communication equilibria in repeated games with incomplete information. *Math Oper Res* 13:191–231
- Forges F (1990) Universal mechanisms. *Econometrica* 58:1341–1364
- Forges F (1993) Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theor Decis* 35:277–310
- Forges F (2006) Correlated equilibrium in games with incomplete information revisited. *Theor Decis* 61:329–344
- Foster D, Vohra R (1997) Calibrated learning and correlated equilibrium. *Games Econ Behav* 21:40–55
- Gerardi D (2000) Interim pre-play communication. Mimeo, Yale University, New Haven
- Gerardi D (2004) Unmediated communication in games with complete and incomplete information. *J Econ Theory* 114:104–131
- Gerardi D, Myerson R (2007) Sequential equilibria in Bayesian games with communication. *Games Econ Behav* 60:104–134
- Gilboa I, Zemel E (1989) Nash and correlated equilibria: some complexity considerations. *Games Econ Behav* 1:80–93
- Gomez-Canovas S, Hansen P, Jaumard B (1999) Nash Equilibria from the correlated equilibria viewpoint. *Int Game Theory Rev* 1:33–44
- Gossner O (1998) Secure protocols or how communication generates correlation. *J Econ Theory* 83:69–89
- Gossner O (2000) Comparison of information structures. *Games Econ Behav* 30:44–63
- Gossner O, Tomala T (2007) Secret correlation in repeated games with signals. *Math Oper Res* 32:413–424
- Halpern JY (2007) Computer science and game theory. In: Durlauf SN, Blume LE (eds) *The New Palgrave dictionary of economics*, 2nd edn. Palgrave Macmillan. The New Palgrave dictionary of economics online. http://www.dictionaryofeconomics.com/article?id=pde2008_C000566. Accessed 24 May 2008
- Hart S, Schmeidler D (1989) Existence of correlated equilibria. *Math Oper Res* 14:18–25
- Hart S, Mas-Colell A (2000) A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68:1127–1150

- Hart S (2005) Adaptive heuristics. *Econometrica* 73:1401–1430
- Krishna RV (2007) Communication in games of incomplete information: two players. *J Econ Theory* 132:584–592
- Lehrer E (1991) Internal correlation in repeated games. *Int J Game Theory* 19:431–456
- Lehrer E (1992) Correlated equilibria in two-player repeated games with non-observable actions. *Math Oper Res* 17:175–199
- Lehrer E (1996) Mediated talk. *Int J Game Theory* 25:177–188
- Lehrer E, Sorin S (1997) One-shot public mediated talk. *Games Econ Behav* 20:131–148
- Lehrer E, Rosenberg D, Shmaya E (2006) Signaling and mediation in Bayesian games. Mimeo, Tel Aviv University, Tel Aviv
- Milgrom P, Roberts J (1996) Coalition-proofness and correlation with arbitrary communication possibilities. *Games Econ Behav* 17:113–128
- Moreno D, Wooders J (1996) Coalition-proof equilibrium. *Games Econ Behav* 17:80–112
- Myerson R (1982) Optimal coordination mechanisms in generalized principal-agent problems. *J Math Econ* 10:67–81
- Myerson R (1986a) Multistage games with communication. *Econometrica* 54:323–358
- Myerson R (1986b) Acceptable and predominant correlated equilibria. *Int J Game Theory* 15:133–154
- Myerson R (1997) Dual reduction and elementary games. *Games Econ Behav* 21:183–202
- Nash J (1951) Non-cooperative games. *Ann Math* 54:286–295
- Nau RF (1992) Joint coherence in games with incomplete information. *Manag Sci* 38:374–387
- Nau RF, McCardle KF (1990) Coherent behavior in non-cooperative games. *J Econ Theory* 50(2):424–444
- Nau RF, McCardle KF (1991) Arbitrage, rationality and equilibrium. *Theor Decis* 31:199–240
- Nau RF, Gomez-Canovas S, Hansen P (2004) On the geometry of Nash equilibria and correlated equilibria. *Int J Game Theory* 32:443–453
- Papadimitriou CH (2005) Computing correlated equilibria in multiplayer games. In: *Proceedings of the 37th ACM symposium on theory of computing. STOC*, Baltimore, pp 49–56
- Ray I (1996) Coalition-proof correlated equilibrium: a definition. *Games Econ Behav* 17:56–79
- Renault J, Tomala T (2004) Communication equilibrium payoffs in repeated games with imperfect monitoring. *Games Econ Behav* 49:313–344
- Solan E (2001) Characterization of correlated equilibrium in stochastic games. *Int J Game Theory* 30:259–277
- Solan E, Vieille N (2002) Correlated equilibrium in stochastic games. *Game Econ Behav* 38:362–399
- Urbano A, Vila J (2002) Computational complexity and communication: coordination in two-player games. *Econometrica* 70:1893–1927
- Urbano A, Vila J (2004a) Computationally restricted unmediated talk under incomplete information. *J Econ Theory* 23:283–320
- Urbano A, Vila J (2004b) Unmediated communication in repeated games with imperfect monitoring. *Games Econ Behav* 46:143–173
- Vida P (2007) From communication equilibria to correlated equilibria. Mimeo, University of Vienna, Vienna
- Viossat Y (2008) Is having a unique equilibrium robust? *J Math Econ* 44:1152–1160
- Viossat Y (2006) The geometry of Nash equilibria and correlated equilibria and a generalization of zero-sum games. Mimeo, S-WoPEc working paper 641. Stockholm School of Economics, Stockholm
- Viossat Y (2007) The replicator dynamics does not lead to correlated equilibria. *Games Econ Behav* 59:397–407

Books and Reviews

- Forges F (1994) Non-zero sum repeated games and information transmission. In: Megiddo N (ed) *Essays in game theory in honor of Michael Maschler*. Springer, Berlin, pp 65–95
- Mertens JF (1994) Correlated- and communication equilibria. In: Mertens JF, Sorin S (eds) *Game theoretic methods in general equilibrium analysis*. Kluwer, Dordrecht, pp 243–248
- Myerson R (1985) Bayesian equilibrium and incentive compatibility. In: Hurwicz L, Schmeidler D, Sonnenschein H (eds) *Social goals and social organization*. Cambridge University Press, Cambridge, pp 229–259
- Myerson R (1994) Communication, correlated equilibria and incentive compatibility. In: Aumann R, Hart S (eds) *Handbook of game theory*, vol 2. Elsevier, Amsterdam, pp 827–847
- Sorin S (1997) Communication, correlation and cooperation. In: Mas Colell A, Hart S (eds) *Cooperation: game theoretic approaches*. Springer, Berlin, pp 198–218

Bayesian Games: Games with Incomplete Information

Shmuel Zamir

Center for the Study of Rationality, Hebrew University, Jerusalem, Israel

Article Outline

Definition

Introduction

Harsanyi's Model: The Notion of *Type*

Aumann's Model

Harsanyi's Model and Hierarchies of Beliefs

The Universal Belief Space

Belief Subspaces

Consistent Beliefs and Common Priors

Bayesian Games and Bayesian Equilibrium

Bayesian Equilibrium and Correlated Equilibrium

Concluding Remarks and Future Directions

Bibliography

Glossary

Bayesian equilibrium A Nash equilibrium of a Bayesian game: A list of behavior and beliefs such that each player is doing his best to maximize his payoff, according to his beliefs about the behavior of the other players.

Bayesian game An interactive decision situation involving several decision makers (players) in which each player has beliefs about (i.e., assigns probability distribution to) the payoff relevant parameters and the beliefs of the other players.

Common prior and consistent beliefs The beliefs of players in a game with incomplete information are said to be consistent if they are derived from the same probability distribution (the common prior) by conditioning on each player's private information. In other words, if the beliefs are consistent, the only source of differences in beliefs is difference in information.

Correlated equilibrium A Nash equilibrium in an extension of the game in which there is a chance move, and each player has only partial information about its outcome.

State of nature Payoff relevant data of the game such as payoff functions, value of a random variable, etc. It is convenient to think of a state of nature as a full description of a "game-form" (actions and payoff functions).

State of the world A specification of the state of nature (payoff relevant parameters) and the players' types (belief of all levels). That is, a state of the world is a state of nature and a list of the states of mind of all players.

Type Also known as *state of mind* and is a full description of player's beliefs (about the state of nature), beliefs about beliefs of the other players, beliefs about the beliefs about his beliefs, etc. ad infinitum.

Definition

Bayesian games (also known as *Games with Incomplete Information*) are models of interactive decision situations in which the decision makers (players) have only partial information about the data of the game and about the other players. Clearly this is typically the situation we are facing and hence the importance of the subject: The basic underlying assumption of classical game theory according to which the data of the game is *common knowledge* (CK) among the players is too strong and often implausible in real situations. The importance of Bayesian games is in providing the tools and methodology to relax this implausible assumption, to enable modeling of the overwhelming majority of real-life situations in which players have only partial information about the payoff relevant data. As a result of the interactive nature of the situation, this methodology turns out to be rather deep and sophisticated, both conceptually and mathematically. Adopting the classical Bayesian approach of

statistics, we encounter the need to deal with an *infinite hierarchy of beliefs*: what does each player believe that the other player believes about what he believes. . . is the actual payoff associated with a certain outcome? It is not surprising that this methodological difficulty was a major obstacle in the development of the theory, and this article is largely devoted to explaining and resolving this methodological difficulty.

Introduction

A game is a mathematical model for an interactive decision situation involving several decision makers (players) whose decisions affect each other. A basic, often implicit, assumption is that the data of the game, which we call the *state of nature*, are *common knowledge* (CK) among the players. In particular the actions available to the players and the payoff functions are CK. This is a rather strong assumption that says that every player knows all actions and payoff functions of all players, every player knows that all other players know all actions and payoff functions, every player knows that every player knows that every player knows, etc. ad infinitum. *Bayesian games* (also known as *games with incomplete information*), which is the subject of this article, are models of interactive decision situations in which each player has only partial information about the payoff relevant parameters of the given situation.

Adopting the Bayesian approach, we assume that a player who has only partial knowledge about the state of nature has some *beliefs*, namely, prior distribution, about the parameters which he does not know or he is uncertain about. However, unlike in a statistical problem which involves a single decision maker, this is not enough in an interactive situation: As the decisions of other players are relevant, so are their beliefs, since they affect their decisions. Thus, a player must have beliefs about the beliefs of other players. For the same reason, a player needs beliefs about the beliefs of other players about his beliefs and so on. This interactive reasoning about beliefs leads unavoidably to *infinite hierarchies of beliefs*

which looks rather intractable. The natural emergence of hierarchies of beliefs is illustrated in the following example:

Example 1 Two players, P1 and P2, play a 2×2 game whose payoffs depend on an unknown state of nature $s \in \{1, 2\}$. Player P1's actions are $\{T, B\}$, player P2's actions are $\{L, R\}$, and the payoffs are given in the following matrices:

		P2	
		L	R
P1	T	0, 1	1, 0
	B	1, 0	0, 1

Payoffs when $s = 1$

		P2	
		L	R
P1	T	1, 0	0, 1
	B	0, 1	1, 0

Payoffs when $s = 2$

Assume that the belief (prior) of P1 about the event $\{s = 1\}$ is p and the belief of P2 about the same event is q . The best action of P1 depends both on his prior and on the action of P2 and similarly for the best action of P2. This is given in the following tables:

		P2's action	
		L	R
$p < 0.5$	T		
	B		
$p > 0.5$			

Best reply of P1

		$q < 0.5$	$q > 0.5$
		R	L
P1's action	T		
	B		

Best reply of P2

Now, since the optimal action of P1 depends not only on his belief p but also on the, unknown to him, action of P2, which depends on his belief q , player P1 must therefore have beliefs about q . These are his *second-level beliefs*, namely, beliefs about beliefs.

But then, since this is relevant and unknown to P2, he must have beliefs about that which will be *third-level beliefs* of P2 and so on. The whole infinite hierarchies of beliefs of the two players pop out naturally in the analysis of this simple two-person game of incomplete information.

The objective of this article is to model this kind of situation. Most of the effort will be devoted to the modeling of the mutual beliefs structure, and only then we add the underlying game which, together with the beliefs structure, defines a *Bayesian game* for which we define the notion of *Bayesian equilibrium*.

Harsanyi's Model: The Notion of Type

As suggested by our introductory example, the straightforward way to describe the mutual beliefs structure in a situation of incomplete information is to specify explicitly the whole hierarchies of beliefs of the players, that is, the beliefs of each player about the unknown parameters of the game, each player's beliefs about the other players' beliefs about these parameters, each player's beliefs about the other players' beliefs about his beliefs about the parameters, and so on ad infinitum. This may be called the *explicit* approach and is in fact feasible and was explored and developed at a later stage of the theory (see Aumann 1999a, b; Aumann and Heifetz 2002; Mertens and Zamir 1985). We will come back to it when we discuss the universal belief space. However, for obvious reasons, the explicit approach is mathematically rather cumbersome and hardly manageable. Indeed this was a major obstacle to the development of the theory of games with incomplete information at its early stages. The breakthrough was provided by John Harsanyi (1967) in a seminal work that earned him the Nobel Prize some 30 years later. While Harsanyi actually formulated the problem verbally, in an explicit way, he suggested a solution that "avoided" the difficulty of having to deal with infinite hierarchies of beliefs, by providing a much more workable *implicit*, encapsulated model which we present now.

The key notion in Harsanyi's model is that of *type*. Each player can be of several types where a type is to be thought of as a full description of the

player's beliefs about the *state of nature* (the data of the game), beliefs about the beliefs of other players about the state of nature and about his own beliefs, etc. One may think of a player's type as his *state of mind*: a specific configuration of his brain that contains an answer to any question regarding beliefs about the state of nature and about the types of the other players. Note that this implies self-reference (of a type to itself through the types of other players) which is unavoidable in an interactive decision situation. A Harsanyi game of incomplete information consists of the following ingredients (to simplify notations, assume all sets to be finite):

- I – Player's set.
- S – The set of states of nature.
- T_i – The type set of player $i \in I$. Let $T = \times_{i \in I} T_i$ denote the *type set*, that is, the set type profiles.
- $Y \subset S \times T$ – A set of *states of the world*.
- $p \in \Delta(Y)$ – Probability distribution on Y , called the *common prior*.

(For a set A , we denote the set of probability distributions on A by $\Delta(A)$).

Remark A state of the world ω thus consists of a state of nature and a list of the types of the players. We denote it as

$$\omega = (s(\omega); t_1(\omega), \dots, t_n(\omega)).$$

We think of the state of nature as a full description of the game which we call a *game-form*. So, if it is a game in strategic form, we write the state of nature at state of the world ω as

$$s(\omega) = \left(I, (A_i(\omega))_{i \in I}, (u_i(\cdot; \omega))_{i \in I} \right).$$

The payoff functions u_i depend only on the state of nature and not on the types. That is, for all $i \in I$,

$$s(\omega) = s(\omega') \Rightarrow u_i(\cdot; \omega) = u_i(\cdot; \omega').$$

The game with incomplete information is played as follows:

1. A chance move chooses $\omega = (s(\omega); t_1(\omega), \dots, t_n(\omega)) \in Y$ using the probability distribution p .

2. Each player is told his chosen type $t_i(\omega)$ (but not the chosen state of nature $s(\omega)$ and not the other players' types $t_{-i}(\omega) = (t_j(\omega))_{j \neq i}$).
3. The players choose simultaneously an action: player i chooses $a_i \in A_i(\omega)$ and receives a payoff $u_i(a; \omega)$ where $a = (a_1, \dots, a_n)$ is the vector of chosen actions and ω is the state of the world chosen by the chance move.

Remark The set $A_i(\omega)$ of actions available to player i in state of the world ω must be known to him. Since his only information is his type $t_i(\omega)$, we must impose that $A_i(\omega)$ is T_i measurable, i.e.,

$$t_i(\omega) = t_i(\omega') \Rightarrow A_i(\omega) = A_i(\omega').$$

Note that if $s(\omega)$ was commonly known among the players, it would be a regular game in strategic form. We use the term “game-form” to indicate that the players have only partial information about $s(\omega)$. The players do not know which $s(\omega)$ is being played. In other words, in the extensive form game of Harsanyi, the game-forms $(s(\omega))_{\omega \in Y}$ are not subgames since they are interconnected by information sets: Player i does not know which $s(\omega)$ is being played since he does not know ω ; he knows only his own type $t_i(\omega)$.

An important application of Harsanyi's model is made in *auction theory*, as an auction is a clear situation of incomplete information. For example, in a closed private-value auction of a single indivisible object, the type of a player is his private value for the object, which is typically known to him and not to other players. We come back to this in the section entitled “[Examples of Bayesian Equilibria](#).”

Aumann's Model

A frequently used model of incomplete information was given by Aumann (1976).

Definition 2 An Aumann model of incomplete information is $(I, Y, (\pi_i)_{i \in I}, P)$ where:

- I is the players' set.
- Y is a (finite) set whose elements are called *states of the world*.

- For $i \in I$, π_i is a partition of Y .
- P is a probability distribution on Y , also called the *common prior*.

In this model, a state of the world $\omega \in Y$ is chosen according to the probability distribution P , and each player i is informed of $\pi_i(\omega)$, the element of his partition that contains the chosen state of the world ω . This is the informational structure which becomes a game with incomplete information if we add a mapping $s: Y \rightarrow S$. The state of nature $s(\omega)$ is the game-form corresponding to the state of the world ω (with the requirement that the action sets $A_i(\omega)$ are π_i measurable).

It is readily seen that Aumann's model is a Harsanyi model in which the type set T_i of player i is the set of his partition elements, i.e., $T_i = \{\pi_i(-\omega) | \omega \in Y\}$, and the common prior on Y is P . Conversely, any Harsanyi model is an Aumann model in which the partitions are those defined by the types, i.e., $\pi_i(\omega) = \{\omega' \in Y | t_i(\omega') = t_i(\omega)\}$.

Harsanyi's Model and Hierarchies of Beliefs

As our starting point in modeling incomplete information situations was the appearance of hierarchies of beliefs, one may ask how is the Harsanyi (or Aumann) model related to hierarchies of beliefs and how does it capture this unavoidable feature of incomplete information situations? The main observation towards answering this question is the following:

Proposition 3 Any state of the world in Aumann's model or any type profile $t \in T$ in Harsanyi's model defines (uniquely) a hierarchy of mutual beliefs among the players.

Let us illustrate the idea of the proof by the following example:


Example Consider a Harsanyi model with two players, I and II , each of which can be of two

types: $T_I = \{I_1, I_2\}$, $T_{II} = \{II_1, II_2\}$ and thus $T = \{(I_1, II_1), (I_1, II_2), (I_2, II_1), (I_2, II_2)\}$. The probability p on types is given by

	II_1	II_2
I_1	$\frac{1}{4}$	$\frac{1}{4}$
I_2	$\frac{1}{3}$	$\frac{1}{6}$

Denote the corresponding states of nature by $a = s(I_1 II_2)$, $b = s(I_1 II_1)$, $c = s(I_2 II_1)$, and $d = s(I_2 II_2)$. These are the *states of nature* about which there is incomplete information.

The game in extensive form:



Assume that the state of nature is a . What are the belief hierarchies of the players?

	II_1	II_2
I_1	a $\frac{1}{4}$	b $\frac{1}{4}$
I_2	c $\frac{1}{3}$	d $\frac{1}{6}$

First-level beliefs are obtained by each player from p , by conditioning on his type:

- I_1 : With probability $\frac{1}{2}$ the state is a and with probability $\frac{1}{2}$ the state is b .

- I_2 : With probability $\frac{2}{3}$ the state is c and with probability $\frac{1}{3}$ the state is d .
- II_1 : With probability $\frac{3}{7}$ the state is a and with probability $\frac{4}{7}$ the state is c .
- II_2 : With probability $\frac{3}{5}$ the state is b and with probability $\frac{2}{5}$ the state is d .

Second-level beliefs (using shorthand notation for the above beliefs: $(\frac{1}{2}a^* + \frac{1}{2}b)$, etc.):

- I_1 : With probability $\frac{1}{2}$, player II believes $(\frac{3}{7}a^* + \frac{4}{7}c)$, and with probability $\frac{1}{2}$, player II believes $(\frac{3}{5}b^* + \frac{2}{5}d)$.
- I_2 : With probability $\frac{2}{3}$, player II believes $(\frac{3}{7}a^* + \frac{4}{7}c)$, and with probability $\frac{1}{3}$, player II believes $(\frac{3}{5}b^* + \frac{2}{5}d)$.
- II_1 : With probability $\frac{3}{7}$, player I believes $(\frac{1}{2}a^* + \frac{1}{2}b)$, and with probability $\frac{4}{7}$, player I believes $(\frac{2}{3}c^* + \frac{1}{3}d)$.
- II_2 : With probability $\frac{3}{5}$, player I believes $(\frac{1}{2}a^* + \frac{1}{2}b)$, and with probability $\frac{2}{5}$, player I believes $(\frac{2}{3}c^* + \frac{1}{3}d)$.

Third-level beliefs:

- I_1 : With probability $\frac{1}{2}$, player II believes that: “With probability $\frac{3}{7}$, player I believes $(\frac{1}{2}a^* + \frac{1}{2}b)$ and with probability $\frac{4}{7}$, player I believes $(\frac{2}{3}c^* + \frac{1}{3}d)$.” And with probability $\frac{1}{2}$, player II believes that: “With probability $\frac{3}{5}$, player I believes $(\frac{1}{2}a^* + \frac{1}{2}b)$ and with probability $\frac{2}{5}$, player I believes $(\frac{2}{3}c^* + \frac{1}{3}d)$.”

The idea is very simple and powerful; since each player of a given type has a probability distribution (beliefs) both about the types of the other players and about the set S of states of nature, the hierarchies of beliefs are constructed inductively: If the k th level beliefs (about S) are defined for each type, then the beliefs about types generate the $(k + 1)$ th level of beliefs.

Thus, the compact model of Harsanyi *does* capture the whole hierarchies of beliefs and it is rather tractable. The natural question is whether this model can be used for *all* hierarchies of beliefs. In other words, given any hierarchy of mutual beliefs of a set of players I about a set S of states of nature, can it be represented by a Harsanyi game? This was answered by Mertens and Zamir (1985), who constructed the universal belief space; that is, given a set S of states of nature and a finite set I of players, they looked for the space Ω of *all possible* hierarchies of mutual beliefs about S among the players in I . This construction is outlined in the next section.

The Universal Belief Space

Given a finite set of players $I = \{1, \dots, n\}$ and a set S of states of nature, which are assumed to be compact, we first identify the mathematical spaces in which lie the hierarchies of beliefs. Recall that $\Delta(A)$ denotes the set of probability distributions on A and defines inductively the sequence of spaces $(X_k)_{k=1}^\infty$ by

$$X_1 = \Delta(S) \quad (1)$$

$$X_{k+1} = X_k \times \Delta(S \times X_k^{n-1}), \text{ for } k = 1, 2, \dots \quad (2)$$

Any probability distribution on S can be a first-level belief and is thus in X_1 . A second-level belief is a joint probability distribution on S and the first-level beliefs of the other $(n - 1)$ players. This is an element in $\Delta(S \times X_1^{n-1})$, and therefore, a two-level hierarchy is an element of the product space $X_1 \times \Delta(S \times X_1^{n-1})$ and so on for any level. Note that at each level belief is a joint probability distribution on S and the previous level beliefs, allowing for correlation between the two. In dealing with these probability spaces, we need to have some mathematical structure. More specifically, we make use of the *weak** topology.

Definition 4 A sequence $(F_n)_{n=1}^\infty$ of probability measures (on Ω) converges in the *weak**

topology to the probability F if and only if $\lim_{n \rightarrow \infty} \int_{\Omega} g(\omega) dF_n = \int_{\Omega} g(\omega) dF$ for all bounded and continuous functions $g : \Omega \rightarrow \mathbb{R}$.

It follows from the compactness of S that all spaces defined by Eqs. 1 and 2 are compact in the *weak** topology. However, for $k > 1$, not every element of X_k represents a coherent hierarchy of beliefs of level k . For example, if $(\mu_1, \mu_2) \in X_2$ where $\mu_1 \in \Delta(S) = X_1$ and $\mu_2 \in \Delta(S \times X_1^{n-1})$, then for this to describe meaningful beliefs of a player, the marginal distribution of μ_2 on S must coincide with μ_1 . More generally, any event A in the space of k -level beliefs has to have the same (marginal) probability in any higher-level beliefs. Furthermore, not only are each player's beliefs coherent, but he also considers only coherent beliefs of the other players (only those that are in support of his beliefs). Expressing formally this coherency condition yields a selection $T_k \subseteq X_k$ such that $T_1 = X_1 = \Delta(S)$. It is proved that the projection of T_{k+1} on X_k is T_k (i.e., any coherent k -level hierarchy can be extended to a coherent $k + 1$ -level hierarchy) and that all the sets T_k are compact. Therefore, the *projective limit*, $T = \lim_{\leftarrow k} T_k$ is well defined and nonempty such that $\mu_{k+1} = (\mu_k, \nu_k)$. (The projective limit (also known as the *inverse limit*) of the sequence $(T_k)_{k=1}^\infty$ is the space T of all sequences $(\mu_1, \mu_2, \dots) \in \prod_{k=1}^\infty T_k$ which satisfy: For any $k \in \mathbb{N}$, there is a probability distribution $\nu_k \in \Delta(S \times T_k^{n-1})$).

Definition 5 The *universal type space* T is the projective limit of the spaces $(T_k)_{k=1}^\infty$.

That is, T is the set of *all coherent infinite hierarchies of beliefs* regarding S , of a player in I . It does not depend on i since by construction it contains *all possible* hierarchies of beliefs regarding S , and it is therefore the same for all players. It is determined only by S and the number of players n .

Proposition 6 The universal type space T is compact and satisfies

$$T \approx \Delta(S \times T^{n-1}). \quad (3)$$

The $\approx$ sign in Eq. 3 is to be read as an *isomorphism* and Proposition 6 says that a type of player

can be identified with a joint probability distribution on the state of nature and the types of the other players. The implicit Eq. 3 reflects the self-reference and circularity of the notion of type: The type of a player is his beliefs about the state of nature and about all the beliefs of the other players, in particular, their beliefs about his own beliefs.

Definition 7 The *universal belief space* (UBS) is the space Ω defined by

$$\Omega = S \times T^n \quad (4)$$

An element of Ω is called a *state of the world*.

Thus, a state of the world is $\omega = (s(\omega); t_1(\omega), t_2(\omega), \dots, t_n(\omega))$ with $s(\omega) \in S$ and $t_i(\omega) \in T$ for all i in I . This is the specification of the states of nature and the types of all players. The universal belief space Ω is what we looked for: the set of *all* incomplete information and mutual belief configurations of a set of n players regarding the state of nature. In particular, as we will see later, all Harsanyi and Aumann models are embedded in Ω , but it includes also belief configurations that cannot be modeled as Harsanyi games. As we noted before, the UBS is determined only by the set of states of nature S and the set of players I , so it should be denoted as $\Omega(S, I)$. For the sake of simplicity, we shall omit the arguments and write Ω , unless we wish to emphasize the underlying sets S and I .

The execution of the construction of the UBS according to the outline above involves some non-trivial mathematics, as can be seen in Mertens and Zamir (1985). The reason is that even with a finite number of states of nature, the space of first-level beliefs is a continuum, the second level is the space of probability distributions on a continuum, and the third level is the space of probability distributions on the space of probability distributions on a continuum. This requires some structure for these spaces: For a (Borel) measurable event E , let $B_i^p(E)$ be the event “player i of type t_i believes that the probability of E is at least p , that is,

$$B_i^p(E) = \{\omega \in \Omega | t_i(E) \geq p\}$$

Since this is the object of beliefs of players other than i (beliefs of $j \neq i$ about the beliefs

of i), this set must also be measurable. Mertens and Zamir used the *weak** topology which is the minimal topology with which the event $B_i^p(E)$ is (Borel) measurable for any (Borel) measurable event E . In this topology, if A is a compact set, then $\Delta(A)$, the space of all probability distributions on A , is also compact. However, the hierarchic construction can also be made with stronger topologies on $\Delta(A)$ (see Brandenburger and Dekel 1993; Heifetz 1993; Mertens et al. 1994). Heifetz and Samet (1998) worked out the construction of the universal belief space without topology, using only a measurable structure (which is implied by the assumption that the beliefs of the players are measurable). All these explicit constructions of the belief space are within what is called the *semantic* approach. Aumann (1999b) provided another construction of a belief system using the *syntactic* approach based on *sentences and logical formulas* specifying explicitly what each player believes about the state of nature, about the beliefs of the other players about the state of nature, and so on. For a detailed construction, see Aumann (1999b), Heifetz and Mongin (2001), and Meier (2001). For a comparison of the syntactic and semantic approaches, see Aumann and Heifetz (2002).

Belief Subspaces

In constructing the universal belief space, we implicitly assumed that each player knows his own type since we specified only his beliefs about the state of nature and about the beliefs of the *other* players. In view of that, and since by Eq. 3 a type of player i is a probability distribution on $S \times T^{\Lambda \setminus \{i\}}$, we can view a type t_i also as a probability distribution on $\Omega = S \times T^I$ in which the marginal distribution on T_i is a degenerate delta function at t_i ; that is, if $\omega = (s(\omega); t_1(\omega), t_2(\omega), \dots, t_n(\omega))$, then for all i in I ,

$$t_i(\omega) \in \Delta(\Omega) \quad \text{and} \quad t_i(\omega)[t_i = t_i(\omega)] = 1. \quad (5)$$

In particular, it follows that if $\text{Supp}(t_i)$ denotes the support of t_i , then

$$\omega' \in \text{Supp}(t_i(\omega)) \Rightarrow t_i(\omega') = t_i(\omega). \quad (6)$$

Let $P_i(\omega) = \text{Supp}(t_i(\omega)) \subseteq \Omega$. This defines a *possibility correspondence*; at state of the world ω , player i does not consider as possible any point not in $P_i(\omega)$. By Eq. 6,

$$P_i(\omega) \cap P_i(\omega') \neq \emptyset \Rightarrow P_i(\omega) = P_i(\omega').$$

However, unlike in Aumann's model, P_i does not define a partition of Ω since it is possible that $\omega \notin P_i(\omega)$, and hence the union $\cup_{\omega \in \Omega} P_i(\omega)$ may be strictly smaller than Ω (see Example 7). If $\omega \in P_i(\omega) \subseteq Y$ holds for all ω in some subspace $Y \subset \Omega$ then $(P_i(\omega))_{\omega \in Y}$ is a partition of Y .

As we said, the universal belief space includes all possible beliefs and mutual belief structures over the state of nature. However, in a specific situation of incomplete information, it may well be that only part of Ω is relevant for describing the situation. If the state of the world is ω , then clearly all states of the world in $\cup_{i \in I} P_i(\omega)$ are relevant, but this is not all, because if $\omega' \in P_i(\omega)$, then all states in $P_j(\omega')$, for $j \neq i$, are also relevant in the considerations of player i . This observation motivates the following definition:

Definition 8 A *belief subspace* (*BL-subspace*) is a closed subset Y of Ω which satisfies

$$P_i(\omega) \subseteq Y \quad \forall i \in I \text{ and } \forall \omega \in Y. \quad (7)$$

A belief subspace is *minimal* if it has no proper subset which is also a belief subspace. Given $\omega \in \Omega$, the belief subspace at ω , denoted by $Y(\omega)$, is the minimal subspace containing ω .

Since Ω is a *BL-subspace*, $Y(\omega)$ is well defined for all $\omega \in \Omega$. A *BL-subspace* is a closed subset of Ω which is also *closed under beliefs* of the players. In any $\omega \in Y$, it contains all states of the world which are relevant to the situation: If $\omega' \notin Y$, then no player believes that ω' is possible, no player believes that any other player believes that ω' is possible, no player believes that any player believes that any player believes, etc.

Remark 9 The subspace $Y(\omega)$ is meant to be the minimal subspace which is belief closed by all players *at the state* ω . Thus, a natural definition

would be $\tilde{Y}(\omega)$ is the minimal *BL-subspace* containing $P_i(\omega)$ for all $i \in I$. However, if for every player the state ω is not in $P_i(\omega)$, then $\omega \notin \tilde{Y}(\omega)$. Yet, even if it is not in the belief closure of the players, the real state ω is still relevant (at least for the analyst) because it determines the true state of nature; that is, it determines the true payoffs of the game. This is the reason for adding the true state of the world ω , even though “it may not be in the mind of the players.”

It follows from Eqs. 5, 6, and 7 that a *BL-subspace* Y has the following structure:

Proposition 10 A closed subset Y of the universal belief space Ω is a *BL-subspace* if and only if it satisfies the following conditions:

1. For any $\omega = (s(\omega); t_1(\omega), t_2(\omega), \dots, t_n(\omega)) \in Y$ and for all i , the type $t_i(\omega)$ is a probability distribution on Y .
2. For any ω and ω' in Y ,

$$\omega' \in \text{Supp}(t_i(\omega)) \Rightarrow t_i(\omega') = t_i(\omega).$$

In fact condition 1 follows directly from Definition 8, while condition 2 follows from the general property of the UBS expressed in Eq. 6.

Given a *BL-subspace* Y in $\Omega(S, I)$, we denote by T_i the *type set of player* i :

$$T_i = \{t_i(\omega) | \omega \in Y\},$$

and note that unlike in the UBS, in a specific model Y , the type sets are typically *not the same* for all i and the analogue of Eq. 4 is

$$Y \subseteq S \times T_1 \times \dots \times T_n.$$

A *BL-subspace* is a model of incomplete information about the state of nature. As we saw in Harsanyi's model, in any model of incomplete information about a fixed set S of states of nature, involving the same set of players I , a state of the world ω defines (encapsulates) an infinite hierarchy of mutual beliefs of the players I on S . By the universality of the belief space $\Omega(S, I)$, there is $\omega' \in \Omega(S, I)$ with the same hierarchy of beliefs as that of ω . The mapping of each ω to its

corresponding ω' in $\Omega(S, I)$ is called a *belief morphism*, as it preserves the belief structure. Mertens and Zamir (1985) proved that the space $\Omega(S, I)$ is universal in the sense that *any* model Y of incomplete information of the set of players I about the state of nature $s \in S$ can be embedded in $\Omega(S, I)$ via belief morphism $\varphi: Y \rightarrow \Omega(S, I)$ so that $\varphi(Y)$ is a belief subspace in $\Omega(S, I)$. In the following examples, we give the *BL*-subspaces representing some known models.

Examples of Belief Subspaces

Example 1 (A Game with Complete Information) If the state of nature is $s_0 \in S$, then in the universal belief space $\Omega(S, I)$, the game is described by a *BL*-subspace Y consisting of a single state of the world:

$$Y = \{\omega\} \quad \text{where } \omega = (s_0; [1\omega], \dots, [1\omega]).$$

Here $[1\omega]$ is the only possible probability distribution on Y , namely, the trivial distribution supported by ω . In particular, the state of nature s_0 (i.e., the data of the game) is commonly known.

Example 2 (Commonly Known Uncertainty About the State of Nature) Assume that the players' set is $I = \{1, \dots, n\}$ and there are k states of nature representing, say, k possible n -dimensional payoff matrices $G_1, \dots, G_k$. At the beginning of the game, the payoff matrix is chosen by a chance move according to the probability vector $p = (p_1, \dots, p_k)$ which is commonly known by the players, but no player receives any information about the outcome of the chance move. The set of states of nature is $S = \{G_1, \dots, G_k\}$. The situation described above is embedded in the UBS, $\Omega(S, I)$, as the following *BL*-subspace Y consisting of k states of the world (denoting $p \in \Delta(Y)$ by $[p_1\omega_1, \dots, p_k\omega_k]$):


- $Y = \{\omega_1, \dots, \omega_k\}$.
- $\omega_1 = (G_1; [p_1\omega_1, \dots, p_k\omega_k], \dots, [p_1\omega_1, \dots, p_k\omega_k])$.
- $\omega_2 = (G_2; [p_1\omega_1, \dots, p_k\omega_k], \dots, [p_1\omega_1, \dots, p_k\omega_k])$.
- $\dots$
- $\omega_k = (G_k; [p_1\omega_1, \dots, p_k\omega_k], \dots, [p_1\omega_1, \dots, p_k\omega_k])$.

There is a single type, $[p_1\omega_1, \dots, p_k\omega_k]$, which is the same for all players. It should be emphasized that the type is a distribution on Y (and not just on the states of nature), which implies that the beliefs $[p_1G_1, \dots, p_kG_k]$ on the state of nature are commonly known by the players.

Example 3 (Two Players with Incomplete Information on One Side) There are two players, $I = \{I, II\}$, and two possible payoff matrices, $S = \{G_1, G_2\}$. The payoff matrix is chosen at random with $P(s = G_1) = p$, known to both players. The outcome of this chance move is known only to player I . Aumann and Maschler have studied such situations in which the chosen matrix is played repeatedly and the issue is how the informed player strategically uses his information (see Aumann and Maschler (1995) and its references). This situation is presented in the UBS by the following *BL*-subspace:

- $Y = \{\omega_1, \omega_2\}$.
- $\omega_1 = (G_1; [1\omega_1], [p\omega_1, (1-p)\omega_2])$.
- $\omega_2 = (G_2; [1\omega_2], [p\omega_1, (1-p)\omega_2])$.

Player I has two possible types: $I_1 = [1\omega_1]$ when he is informed of G_1 and $I_2 = [1\omega_2]$ when he is informed of G_2 . Player II has only one type, $II = [p\omega_1, (1-p)\omega_2]$. We describe this situation in the following *extensive form-like* figure in which the oval forms describe the types of the players in the various vertices.



Example 4 (Incomplete Information About the Other Players' Information) In the next example, taken from Sorin and Zamir (1985), one of two players always knows the state of nature but

one-to-one mapping between the type set T and the set S of states of nature, the situation is generated by a chance move choosing the state of nature $s_{ij} \in S$ according to the distribution p (i.e., $P(s_{ij}) = P(I_i, II_j)$ for i and j in $\{1, 2\}$), and then player I is informed of i and player II is informed of j . As a matter of fact, *all* the BL -subspaces in the previous examples can also be written as Harsanyi games, mostly in a trivial way.

Example 6 (Inconsistent Beliefs) In the same universal belief space, $\Omega(S, I)$ of the previous example, consider now another BL -subspace $\tilde{Y}$ which differs from Y only by changing the type II_1 of player II from $[\frac{3}{5}\omega_{11}, \frac{2}{5}\omega_{21}]$ to $[\frac{1}{2}\omega_{11}, \frac{1}{2}\omega_{21}]$, that is,

- $\tilde{Y} = \{\omega_{11}, \omega_{12}, \omega_{21}, \omega_{22}\}$
- $\omega_{11} = (s_{11}; [\frac{3}{7}\omega_{11}, \frac{4}{7}\omega_{12}], [\frac{1}{2}\omega_{11}, \frac{1}{2}\omega_{21}])$
- $\omega_{12} = (s_{12}; [\frac{3}{7}\omega_{11}, \frac{4}{7}\omega_{12}], [\frac{4}{5}\omega_{12}, \frac{1}{5}\omega_{22}])$
- $\omega_{21} = (s_{21}; [\frac{2}{3}\omega_{21}, \frac{1}{3}\omega_{22}], [\frac{1}{2}\omega_{11}, \frac{1}{2}\omega_{21}])$
- $\omega_{22} = (s_{22}; [\frac{2}{3}\omega_{21}, \frac{1}{3}\omega_{22}], [\frac{4}{5}\omega_{12}, \frac{1}{5}\omega_{22}])$

with type sets

$$T_I = \{I_1, I_2\} = \left\{ \left[\frac{3}{7}\omega_{11}, \frac{4}{7}\omega_{12} \right], \left[\frac{2}{3}\omega_{21}, \frac{1}{3}\omega_{22} \right] \right\}$$

$$T_{II} = \{II_1, II_2\} = \left\{ \left[\frac{1}{2}\omega_{11}, \frac{1}{2}\omega_{21} \right], \left[\frac{4}{5}\omega_{12}, \frac{1}{5}\omega_{22} \right] \right\}.$$

Now, the mutual beliefs about each other's type are

	II_1	II_2
I_1	$3/7$	$4/7$
I_2	$2/3$	$1/3$

Beliefs of player I

	II_1	II_2
I_1	$1/2$	$4/5$
I_2	$1/2$	$1/5$

Beliefs of player II

Unlike in the previous example, these beliefs *cannot be derived from a prior distribution* p . According to Harsanyi, these are *inconsistent beliefs*. A BL -subspace with inconsistent beliefs *cannot be described* as a Harsanyi or Aumann model; it cannot be described as a game in extensive form.

Example 7 (“Highly Inconsistent” Beliefs) In the previous example, even though the beliefs of the players were inconsistent in all states of the world, the true state was considered possible by all players (e.g., in the state ω_{12} , player I assigns to this state probability $4/7$ and player II assigns to it probability $4/5$). As was emphasized before, the UBS contains *all* belief configurations, including *highly inconsistent* or *wrong* beliefs, as the following example shows. The belief subspace of the two players I and II concerning the state of nature which can be s_1 or s_2 is given by

- $Y = \{\omega_1, \omega_2\}$.
- $\omega_1 = (s_1; [\frac{1}{2}\omega_1, \frac{1}{2}\omega_2] [1\omega_2])$.
- $\omega_2 = (s_2; [\frac{1}{2}\omega_1, \frac{1}{2}\omega_2] [1\omega_2])$.

In the state of the world ω_1 , the state of nature is s_1 , player I assigns equal probabilities to s_1 and s_2 , but player II assigns probability 1 to s_2 . In other words, he does not consider as possible the true state of the world (and also the true state of nature): $\omega_1 \notin P_{II}(\omega_1)$ and consequently $\cup_{\omega \in Y} P_{II}(\omega) = \{\omega_2\}$ which is strictly smaller than Y . By the definition of belief subspace and condition 6, this also implies that $\cup_{\omega \in \Omega} P_{II}(\omega)$ is strictly smaller than Ω (as it does not contain ω_1).

Consistent Beliefs and Common Priors

A BL -subspace Y is a semantic belief system presenting, via the notion of types, the hierarchies of belief of a set of players having incomplete information about the state of nature. A state of the world captures the situation at what is called the *interim stage*: Each player knows his own type and has beliefs about the state of nature and the types of the other players. The question “what is the *real state of the world* ω ?” is not addressed. In a BL -subspace, there is no chance move with explicit probability distribution that chooses the state of the world, while such a probability distribution is part of a Harsanyi or an Aumann model. Yet, in the belief space Y of Example 5 in the previous section, such a *prior distribution* p *emerged endogenously* from the structure of Y .

More specifically, if the state $\omega \in Y$ is chosen by a chance move according to the probability distribution p and each player i is told his type $t_i(\omega)$, then his beliefs are *precisely* those described by $t_i(\omega)$. This is a property of the *BL*-subspace that we call *consistency* (which does not hold, for instance, for the *BL*-subspace $\tilde{Y}$ in Example 6) and that we define now: Let $Y \subseteq \Omega$ be a *BL*-subspace.

Definition 11

1. A probability distribution $p \in \Delta(Y)$ is said to be *consistent* if for any player $i \in I$,

$$p = \int_Y t_i(\omega) dp. \quad (8)$$

2. A *BL*-subspace Y is said to be consistent if there is a consistent probability distribution p with $\text{Supp}(p) = Y$. A consistent *BL*-subspace will be called a *C*-subspace. A state of the world $\omega \in \Omega$ is said to be consistent if it is a point in a *C*-subspace.

The interpretation of Eq. 8 is that the probability distribution p is “the average” of the types $t_i(\omega)$ of player i (which are also probability distributions on Y), when the average is taken on Y according to p . This definition is not transparent; it is not clear how it captures the consistency property we have just explained, in terms of a chance move choosing $\omega \in Y$ according to p . However, it turns out to be equivalent.

For $\omega \in Y$, denote $\pi_i(\omega) = \{\omega' \in Y \mid t_i(\omega') = t_i(\omega)\}$; then we have

Proposition 12 A probability distribution $p \in \Delta(Y)$ is consistent if and only if

$$t_i(\omega)(A) = p(A \mid \pi_i(\omega)) \quad (9)$$

holds for all $i \in I$ and for any measurable set $A \subseteq Y$.

In particular, a Harsanyi or an Aumann model is represented by a consistent *BL*-subspace since, by construction, the beliefs are derived from a common prior distribution

which is part of the data of the model. The role of the prior distribution p in these models is actually not that of an additional parameter of the model but rather that of an additional assumption on the belief system, namely, the *consistency assumption*. In fact, if a minimal belief subspace is consistent, then the common prior p is uniquely determined by the beliefs, as we saw in Example 5; *there is no need to specify p as additional data of the system.*

Proposition 13 If $\omega \in \Omega$ is a consistent state of the world, and if $Y(\omega)$ is the smallest consistent *BL*-subspace containing ω , then the consistent probability distribution p on $Y(\omega)$ is uniquely determined.

(The formulation of this proposition requires some technical qualification if $Y(\omega)$ is a continuum).

The consistency (or the existence of a common prior) is quite a strong assumption. It assumes that differences in beliefs (i.e., in probability assessments) are due *only* to differences in information; players having precisely *the same information* will have precisely *the same beliefs*. It is no surprise that this assumption has strong consequences, the most known of which is due to Aumann (1976): Players with consistent beliefs *cannot agree to disagree*. That is, if at some state of the world it is commonly known that one player assigns probability q_1 to an event E and another player assigns probability q_2 to the same event, then it must be the case that $q_1 = q_2$. Variants of this result appear under the title of “No trade theorems” (see, e.g., Milgrom and Stokey 1982): Rational players with consistent beliefs cannot believe that they both can gain from a trade or a bet between them.

The plausibility and the justification of the common prior assumption were extensively discussed in the literature (see, e.g., Aumann 1998; Gul 1998; Harsanyi 1967). It is sometimes referred to in the literature as the *Harsanyi doctrine*. Here we only make the observation that within the set of *BL*-subspaces in Ω , the set of consistent *BL*-subspaces is a set of measure zero. To see the idea of the proof, consider the following example:

Example 8 (Generalization of Examples 5 and 6) Consider a *BL*-subspace as in Examples 5 and 6 but with type sets:

$$\begin{aligned} T_I &= \{I_1, I_2\} \\ &= \{[\alpha_1 \omega_{11}, (1 - \alpha_1) \omega_{12}], [\alpha_2 \omega_{21}, (1 - \alpha_2) \omega_{22}]\} \\ T_{II} &= \{II_1, II_2\} \\ &= \{[\beta_1 \omega_{11}, (1 - \beta_1) \omega_{21}], [\beta_2 \omega_{12}, (1 - \beta_2) \omega_{22}]\}. \end{aligned}$$

For any $(\alpha_1, \alpha_2, \beta_1, \beta_2) \in [0, 1]^4$, this is a *BL*-subspace. The mutual beliefs about each other's type are

	II_1	II_2		II_1	II_2
I_1	α_1	$1 - \alpha_1$	I_1	β_1	β_2
I_2	α_2	$1 - \alpha_2$	I_2	$1 - \beta_1$	$1 - \beta_2$

Beliefs of player *I* Beliefs of player *II*

If the subspace is consistent, these beliefs are obtained as conditional distributions from some prior probability distribution p on $T = T_I \times T_{II}$, say, by p of the following matrix:

	II_1	II_2
I_1	p_{11}	p_{12}
I_2	p_{21}	p_{22}

Prior distribution p on T

This implies (assuming $p_{ij} \neq 0$ for all i and j)

$$\begin{aligned} \frac{p_{11}}{p_{12}} &= \frac{\alpha_1}{1 - \alpha_1}; \quad \frac{p_{21}}{p_{22}} = \frac{\alpha_2}{1 - \alpha_2} \\ \text{and hence } \frac{p_{11}p_{22}}{p_{12}p_{21}} &= \frac{\alpha_1}{1 - \alpha_1} \frac{1 - \alpha_2}{\alpha_2}. \end{aligned}$$

Similarly,

$$\begin{aligned} \frac{p_{11}}{p_{21}} &= \frac{\beta_1}{1 - \beta_1}; \quad \frac{p_{12}}{p_{22}} = \frac{\beta_2}{1 - \beta_2} \\ \text{and hence } \frac{p_{11}p_{22}}{p_{12}p_{21}} &= \frac{\beta_1}{1 - \beta_1} \frac{1 - \beta_2}{\beta_2}. \end{aligned}$$

It follows that the types must satisfy

$$\frac{\alpha_1}{1 - \alpha_1} \frac{1 - \alpha_2}{\alpha_2} = \frac{\beta_1}{1 - \beta_1} \frac{1 - \beta_2}{\beta_2}, \quad (10)$$

which is *generally* not the case. More precisely, the set of $(\alpha_1, \alpha_2, \beta_1, \beta_2) \in [0, 1]^4$ satisfying the condition 10 is a set of measure zero; it is a three-dimensional set in the four-dimensional set $[0, 1]^4$. Nyarko (1991) proved that even the ratio of the dimensions of the set of consistent *BL*-subspaces to the dimension of the set of *BL*-subspaces goes to zero as the latter goes to infinity. Summing up, *most BL*-subspaces are inconsistent and thus do not satisfy the common prior condition.

Bayesian Games and Bayesian Equilibrium

As we said, a game with incomplete information played by Bayesian players, often called a *Bayesian game*, is a game in which the players have incomplete information about the data of the game. Being a Bayesian, each player has beliefs (probability distribution) about any relevant data he does not know, including the beliefs of the other players. So far, we have developed the belief structure of such a situation which is a *BL*-subspace Y in the universal belief space $\Omega(S, I)$. Now we add the action sets and the payoff functions. These are actually part of the description of the state of nature: The mapping $s: \Omega \rightarrow S$ assigns to each state of the world ω the *game-form* $s(\omega)$ played at this state. To emphasize this interpretation of $s(\omega)$ as a game-form, we denote it also as Γ_ω :

$$\Gamma_\omega = (I, A_i(t_i(\omega))_{i \in I}, (u_i(\omega))_{i \in I}),$$

where $A_i(t_i(\omega))$ is the actions set (pure strategies) of player i at ω and $u_i(\omega) : A(\omega) \rightarrow \mathbb{R}$ is his payoff function and $A(\omega) = \times_{i \in I} A_i(t_i(\omega))$ is the set of action profiles at state ω . Note that while the actions of a player depend only on his type, his payoff depends on the actions and types of all the players. For a vector of actions $a \in A(\omega)$, we write $u_i(\omega; a)$ for $u_i(\omega)(a)$. Given a *BL*-subspace $Y \subseteq \Omega(S, I)$, we define the *Bayesian game on Y* as follows:

Definition 14 The Bayesian game on Y is a vector payoff game in which:

- $I = \{1, \dots, n\}$ – the players' set.
- Σ_i – the strategy set of player i is the set of mappings.

$\sigma_i : Y \rightarrow A_i$ which are T_i – measurable.

- In particular,

$$t_i(\omega_1) = t_i(\omega_2) \Rightarrow \sigma_i(\omega_1) = \sigma_i(\omega_2).$$

- Let $\Sigma = \times_{i \in I} \Sigma_i$.
- The payoff function u_i for player i is a *vector-valued function* $u_i = (u_{ti})_{t_i \in T_i}$, where u_{ti} (the payoff function of player i of type t_i) is a mapping

$$u_{ti} : \Sigma_{-i} \rightarrow \mathbb{R}$$

- Defined by

$$u_{ti}(\sigma) = \int_Y u_i(\omega; \sigma(\omega)) dt_i(\omega). \quad (11)$$

Note that u_{ti} is T_i measurable, as it should be. When Y is a finite *BL*-subspace, the above-defined Bayesian game is an n -person “game” in which the payoff for player i is a vector with a payoff for each one of his types (therefore, a vector of dimension $|T_i|$). It becomes a *regular game-form* for a given state of the world ω since then the payoff to player i is $u_{ti(\omega)}$. However, these game-forms are not regular games since they are *interconnected*; the players do not know which of these “games” they are playing (since they do not know the state of the world ω). Thus, just like a Harsanyi game, a Bayesian game on a *BL*-subspace Y consists of a family of *connected* game-forms, one for each $\omega \in Y$. However, unlike a Harsanyi game, a Bayesian game has no chance move that chooses the state of the world (or the vector of types). A way to transform a Bayesian game into a regular game was suggested by R. Selten and was named by Harsanyi as the *Selten game* G^{**} (see p. 496 in (Harsanyi 1967)). This is a game with $|T_1| \cdot |T_2| \cdot \dots \cdot |T_n|$ players (one for each type) in which each player $t_i \in T_i$ chooses a strategy and then *selects* his $(n - 1)$ partners, one from each T_j ; $j \neq i$, according to his beliefs t_i .

Bayesian Equilibrium

Although a Bayesian game is not a regular game, the Nash equilibrium concept based on the notion

of *best reply* can be adapted to yield the solution concept of *Bayesian equilibrium* (also called *Nash-Bayes equilibrium*).

Definition 15 A vector of strategies $\sigma = (\sigma_1, \dots, \sigma_n)$, in a Bayesian game, is called a *Bayesian equilibrium* if for all i in I and for all t_i in T_i ,

$$u_{ti}(\sigma) \geq u_{ti}(\sigma_{-i}; \tilde{\sigma}_i), \quad \forall \tilde{\sigma}_i \in \Sigma_i, \quad (12)$$

where, as usual, $\sigma_{-i} = (\sigma_j)_{j \neq i}$ denotes the vector of strategies of players other than i .

Thus, a Bayesian equilibrium specifies a behavior for each player which is a *best reply* to what he believes is the behavior of the other players, that is, a best reply to the strategies of the other players given his type. In a game with complete information, which corresponds to a *BL*-subspace with one state of the world ($Y = \{\omega\}$), as there is only one type of each player, and the beliefs are all probability one on a singleton, the Bayesian equilibrium is just the well-known Nash equilibrium.

Remark 16 It is readily seen that when Y is finite, any Bayesian equilibrium is a Nash equilibrium of the Selten game G^{**} in which each type is a player who selects the types of his partners according to his beliefs. Similarly, we can transform the Bayesian game into an ordinary game in strategic form by defining the payoff function to player i to be $\tilde{u}_i \sum_{t_i \in T_i} \gamma_{t_i} u_{ti}$ where γ_{t_i} are strictly positive. Again, independently of the values of the constants γ_{t_i} , any Bayesian equilibrium is a Nash equilibrium of this game and vice versa. In particular, if we choose the constants so that $\sum_{t_i \in T_i} \gamma_{t_i} = 1$, we obtain the game suggested by Aumann and Maschler in 1967 (see p. 95 in Aumann and Maschler 1995), and again, the set of Nash equilibria of this game is precisely the set of Bayesian equilibria.

The Harsanyi Game Revisited

As we observed in Example 5, the belief structure of a consistent *BL*-subspace is the same as in a Harsanyi game *after the chance move choosing the types*. That is, the embedding of the Harsanyi

game as a *BL*-subspace in the universal belief space is only at *the interim stage, after* the moment that each player gets to know his type. The Harsanyi game on the other hand is at the *ex ante stage, before* a player knows his type. Then, what is the relation between the Nash equilibrium in the Harsanyi game at the *ex ante* stage and the equilibrium at the interim stage, namely, the Bayesian equilibrium of the corresponding *BL*-subspace? This is an important question concerning the embedding of the Harsanyi game in the UBS since, as we said before, the chance move choosing the types *does not appear explicitly in the UBS*. The answer to this question was given by Harsanyi (1967–1968) (assuming that each type t_i has a positive probability).

Theorem 17 (Harsanyi) The set of Nash equilibria of a Harsanyi game is identical to the set of Bayesian equilibria of the equivalent *BL*-subspace in the UBS.

In other words, this theorem states that any equilibrium in the *ex ante* stage is also an equilibrium at the interim stage and vice versa.

In modeling situations of incomplete information, the interim stage is the natural one; if a player *knows his beliefs (type)*, then why should he analyze the situation, as Harsanyi suggests, from the *ex ante* point of view as if his type was not known to him and he could equally well be of another type? Theorem 17 provides a technical answer to this question: The equilibria are the *same* in both games and the equilibrium strategy of the *ex ante* game specifies for each type precisely his equilibrium strategy at the interim stage. In that respect, for a player who knows his type, the Harsanyi model is just an *auxiliary game* to compute his equilibrium behavior. Of course the deeper answer to the question above comes from the interactive nature of the situation: Even though player i knows he is of type t_i , he knows that his partners do not know that and that they may consider the possibility that he is of type $\tilde{t}_i$, and since this affects their behavior, the behavior of type $\tilde{t}_i$ is also relevant to player i who knows he is of type t_i . Finally, Theorem 17 makes the

Bayesian equilibrium the natural extension of the Nash equilibrium concept to games with incomplete information for consistent or inconsistent beliefs, when the Harsanyi ordinary game model is unavailable.

Examples of Bayesian Equilibria

In Example 6, there are two players of two types each and with *inconsistent* mutual beliefs given by

	II_1	II_2
I_1	3/7	4/7
I_2	2/3	1/3

Beliefs of player I

	II_1	II_2
I_1	1/2	4/5
I_2	1/2	1/5

Beliefs of player II

Assume that the payoff matrices for the four types of profiles are

		II	
		L	R
I	T	2, 0	0, 1
	B	0, 0	1, 0

G_{11} : Payoffs when $t = (I_1, II_1)$

		II	
		L	R
I	T	0, 0	0, 0
	B	1, 1	1, 0

G_{12} : Payoffs when $t = (I_1, II_2)$

		II	
		L	R
I	T	0, 0	0, 0
	B	1, 1	0, 0

G_{21} : Payoffs when $t = (I_2, II_1)$

		II	
		L	R
I	T	0, 0	2, 1
	B	0, 0	0, 2

G_{22} : Payoffs when $t = (I_2, II_2)$

As the beliefs are inconsistent, they cannot be presented by a Harsanyi game. Yet we can compute the Bayesian equilibrium of this Bayesian game. Let (x, y) be the strategy of player I , which is:

- Play the mixed strategy $[x(T), (1 - x)(B)]$ when you are of type I_1 .
- Play the mixed strategy $[y(T), (1 - y)(B)]$ when you are of type I_2 .

and let (z, t) be the strategy of player II , which is:

- Play the mixed strategy $[z(L), (1 - z)(R)]$ when you are of type II_1 .
- Play the mixed strategy $[t(L), (1 - t)(R)]$ when you are of type II_2 .

For $0 < x, y, z, t < 1$, each player of each type must be indifferent between his two pure actions; that yields the values in equilibrium:

$$x = \frac{3}{5}, y = \frac{2}{5}, z = \frac{7}{9}, t = \frac{2}{9}.$$

There is no “expected payoff” since this is a Bayesian game and not a game; the expected payoffs depend on the *actual* state of the world, i.e., the actual types of the players and the actual payoff matrix. For example, the state of the world is $\omega_{11} = (G_{11}, I_1, II_1)$; the expected payoffs are

$$\pi(\omega_{11}) = \left(\frac{3}{5}, \frac{2}{5}\right) G_{11} \left(\frac{7/9}{2/9}\right) = \left(\frac{46}{45}, \frac{6}{45}\right).$$

Similarly,

$$\pi(\omega_{12}) = \left(\frac{3}{5}, \frac{2}{5}\right) G_{12} \left(\frac{2/9}{7/9}\right) = \left(\frac{18}{45}, \frac{4}{45}\right)$$

$$\pi(\omega_{21}) = \left(\frac{2}{5}, \frac{3}{5}\right) G_{21} \left(\frac{7/9}{2/9}\right) = \left(\frac{21}{45}, \frac{21}{45}\right)$$

$$\pi(\omega_{22}) = \left(\frac{2}{5}, \frac{3}{5}\right) G_{22} \left(\frac{2/9}{7/9}\right) = \left(\frac{28}{45}, \frac{70}{45}\right).$$

However, these are the *objective* payoffs as viewed by the analyst; they are viewed differently by the players. For player i of type t_i , the relevant payoff is his subjective payoff $u_{it}(\sigma)$ defined in Eq. 11. For example, at state ω_{11} (or ω_{12}), player I believes that with probability $3/7$ the state is ω_{11} in which case his payoff is $46/45$ and with probability $4/7$ the state is ω_{12} in which case his payoff is $18/45$. Therefore, his *subjective* expected payoff at state ω_{11} is $3/7 \times 46/45 + 4/7 \times 18/45 = 2/3$. Similar computations show that in states ω_{21} or ω_{22} , player I “expects” a payoff of $7/15$, while player II “expects” $3/10$ in state ω_{11} or ω_{21} and $86/225$ in state ω_{12} or ω_{22} .

Bayesian equilibrium is widely used in auction theory, which constitutes an important and successful application of the theory of games with incomplete information. The simplest example is that of two buyers bidding in a first-price auction for an indivisible object. If each buyer i has a private value v_i for the object (which is independent of the private value v_j of the other buyer), and

if he further believes that v_j is random with uniform probability distribution on $[0,1]$, then this is a Bayesian game in which the type of a player is his private valuation; that is, the type sets are $T_1 = T_2 = [0,1]$, which is a continuum. This is a consistent Bayesian game (that is, a Harsanyi game) since the beliefs are derived from the uniform probability distribution on $T_1 \times T_2 = [0,1]^2$. A Bayesian equilibrium of this game is that in which each player bids half of his private value: $b_i(v_i) = v_i/2$ (see, e.g., Chap. III in Wolfstetter (1999)). Although auction theory was developed far beyond this simple example, almost all the models studied so far are Bayesian games with consistent beliefs, that is, Harsanyi games. The main reason of course is that consistent Bayesian games are more manageable since they can be described in terms of an equivalent ordinary game in strategic form. However, inconsistent beliefs are rather plausible and exist in the market place in general and even more so in auction situations. An example of that is the case of collusion of bidders: When a bidding ring is formed, it may well be the case that some of the bidders outside the ring are unaware of its existence and behave under the belief that all bidders are competitive. The members of the ring may or may not know whether the other bidders know about the ring, or they may be uncertain about it. This rather plausible mutual belief situation is typically inconsistent and has to be treated as an inconsistent Bayesian game for which a Bayesian equilibrium is to be found.

Bayesian Equilibrium and Correlated Equilibrium

Correlated equilibrium was introduced in Aumann (1974) as the Nash equilibrium of a game extended by adding to it random events about which the players have partial information. Basically, starting from an ordinary game, Aumann added a probability space and information structure and obtained a game with incomplete information, the equilibrium of which he called a *correlated equilibrium of the original game*. The fact that the Nash equilibrium of a

game with incomplete information is the Bayesian equilibrium suggests that the concept of correlated equilibrium is closely related to that of Bayesian equilibrium. In fact Aumann noticed that and discussed it in a second paper entitled “Correlated equilibrium as an expression of Bayesian rationality” (Aumann 1987). In this section, we review briefly, by way of an example, the concept of correlated equilibrium and state formally its relation to the concept of Bayesian equilibrium.

Example 18 Consider a two-person game with actions $\{T, B\}$ for player 1 and $\{L, R\}$ for player 2 with corresponding payoffs given in the following matrix:

		2	
		L	R
1	T	6, 6	2, 7
	B	7, 2	0, 0

G : Payoffs of the basic game

This game has three Nash equilibria: (T, R) with payoff $(2, 7)$, (B, L) with payoff $(7, 2)$, and the mixed equilibrium $([\frac{2}{3}(T), \frac{1}{3}(B)], [\frac{2}{3}(L), \frac{1}{3}(R)])$ with payoff $(4\frac{2}{3}, 4\frac{2}{3})$. Suppose that we add to the game a chance move that chooses an element in $\{T, B\} \times \{L, R\}$ according to the following probability distribution μ :

		L	R
T		1/3	1/3
	B	1/3	0, 0

μ : Probability distribution on $\{T, B\} \times \{L, R\}$

Let us now extend the game G to a game with incomplete information G^* in which a chance move chooses an element in $\{T, B\} \times \{L, R\}$ according to the probability distribution above. Then, player 1 is informed of the first (left) component of the chosen element and player 2 is informed of the second (right) component. Then, each player chooses an action in G and the payoff is made. If we interpret the partial information as a

suggestion of which action to choose, then it is readily verified that following the suggestion is a Nash equilibrium of the extended game yielding a payoff $(5, 5)$. This was called by Aumann a *correlated equilibrium of the original game G* . In our terminology, the extended game G^* is a Bayesian game and its Nash equilibrium is its Bayesian equilibrium. Thus, what we have here is that a correlated equilibrium of a game is just the Bayesian equilibrium of its extension to a game with incomplete information. We now make this a general formal statement. For simplicity, we use the Aumann model of a game with incomplete information.

Let $G = (I, (A_i)_{i \in I}, (u_i)_{i \in I})$ be a game in strategic form where I is the set of players, A_i is the set of actions (pure strategies) of player i , and u_i is his payoff function.

Definition 19 Given a game in strategic form G , an incomplete information extension (the I -extension) of the game G is the game G^* given by

$$G^* = \left(I, (A_i)_{i \in I}, (u_i)_{i \in I}, (Y, p) \right), (\pi_i)_{i \in I},$$

where (Y, p) is a finite probability space and π_i is a partition of Y (the information partition of player i).

This is an Aumann model of incomplete information, and as we noted before, it is also a Harsanyi-type-based model in which the type of player i at state $\omega \in Y$ is $t_i(\omega) = \pi_i(\omega)$ and a strategy of player i is a mapping from his type set to his mixed actions: $\sigma_i: T_i \rightarrow \Delta(A_i)$.

We identify a correlated equilibrium in the game G by the probability distribution μ on the vectors of actions $A = A_1 \times \dots \times A_n$. Thus, $\mu \in \Delta(A)$ is a correlated equilibrium of the game G if when $a \in A$ is chosen according to μ and each player i is suggested to play a_i , his best reply is in fact to play the action a_i .

Given a game with incomplete information G^* as in definition 19, any vector of strategies of the players $\sigma = (\sigma_1, \dots, \sigma_n)$ induces a probability distribution on the vectors of actions $a \in A$. We denote this as $\mu_\sigma \in \Delta(A)$.

We can now state the relation between correlated and Bayesian equilibria.

Theorem 20 Let σ be a Bayesian equilibrium in the game of incomplete information $G^* = (I, (A_i)_{i \in I}, (u_i)_{i \in I}, (Y, p), (\pi_i)_{i \in I})$; then the induced probability distribution μ_σ is a correlated equilibrium of the basic game $G = (I, (A_i)_{i \in I}, (u_i)_{i \in I})$.

The other direction is

Theorem 21 Let μ be a correlated equilibrium of the game $G = (I, (A_i)_{i \in I}, (u_i)_{i \in I})$; then G has an extension to a game with incomplete information $G^* = (I, (A_i)_{i \in I}, (u_i)_{i \in I}, (Y, p), (\pi_i)_{i \in I})$ with a Bayesian equilibrium σ for which $\mu_\sigma = \mu$.

Concluding Remarks and Future Directions

The Consistency Assumption

To the heated discussion of the merits and justification of the consistency assumption in economic and game-theoretical models, we would like to add a couple of remarks. In our opinion, the appropriate way of modeling an incomplete information situation is at the *interim stage*, that is, when a player *knows his own beliefs (type)*. The Harsanyi *ex ante* model is just an auxiliary construction for the analysis. Actually, this was also the view of Harsanyi, who justified his model by proving that it provides the same equilibria as the interim stage situation it generates (Theorem 17). The *Harsanyi doctrine* says roughly that our models “should be consistent” and if we get an inconsistent model, it must be the case that it not be a “correct” model of the situation at hand. This becomes less convincing if we agree that the interim stage is what we are interested in: Not only are *most* mutual beliefs inconsistent, as we saw in the section entitled “[Consistent Beliefs and Common Priors](#)” above, but it is hard to argue convincingly that the model in Example 5 describes an *adequate* mutual belief situation while the model in Example 6 does not; the only difference between the two is that in one model, a certain type’s beliefs are $[\frac{3}{5}\omega_{11}, \frac{2}{5}\omega_{21}]$, while in the other model his beliefs are $[\frac{1}{2}\omega_{11}, \frac{1}{2}\omega_{21}]$.

Another related point is the fact that if players’ beliefs are the data of the situation (in the interim stage), then these are typically imprecise and rather hard to measure. Therefore, any meaningful result of our analysis should be robust to small changes in the beliefs. This cannot be achieved within the consistent belief systems which are a thin set of measure zero in the universal belief space.

Knowledge and Beliefs

Our interest in this article was mostly in the notion of *beliefs* of players and less in the notion of *knowledge*. These are two related but different notions. Knowledge is defined through a knowledge operator satisfying some axioms. Beliefs are defined by means of probability distributions. Aumann’s model, discussed in the section entitled “[Aumann’s Model](#)” above, has both elements: The knowledge was generated by the partitions of the players, while the beliefs were generated by the probability P on the space Y (and the partitions). Being interested in the subjective beliefs of the player, we could understand “at state of the world $\omega \in \Omega$, player i knows the event $E \subseteq \Omega$ ” to mean “at state of the world $\omega \in \Omega$, player i assigns to the event $E \subseteq \Omega$ probability 1.” However, in the universal belief space, “belief with probability 1” does not satisfy a central axiom of the knowledge operator. Namely, if at $\omega \in \Omega$ player i knows the event $E \subseteq \Omega$, then $\omega \in E$. That is, if a player *knows* an event, then this event in fact happened. In the universal belief space where all coherent beliefs are possible, in a state $\omega \in \Omega$ a player may assign probability 1 to the event $\{\omega'\}$ where $\omega' \neq \omega$. In fact, if in a BL -subspace Y the condition $\omega \in P_i(\omega)$ is satisfied for all i and all $\omega \in Y$, then belief with probability 1 is a knowledge operator on Y . This in fact was the case in Aumann’s and in Harsanyi’s models where, by construction, the support of the beliefs of a player in the state ω always included ω . For a detailed discussion of the relationship between knowledge and beliefs in the universal belief space, see Vassilakis and Zamir (1993).

Future Directions

We have not said much about the existence of Bayesian equilibrium, mainly because it has not been studied enough and there are no general results, especially in the non-consistent case. We can readily see that a Bayesian game on a finite *BL*-subspace in which each state of nature $s(\omega)$ is a finite game-form has a Bayesian equilibrium in mixed strategies. This can be proved, for example, by transforming the Bayesian game into an ordinary finite game (see Remark 16) and applying the Nash theorem for finite games. For games with incomplete information with a continuum of strategies and payoff functions not necessarily continuous, there are no general existence results. Even in consistent auction models, existence was proved for specific models separately (see Maskin and Riley 2000; Milgrom and Weber 1982; Reny and Zamir 2004). Establishing general existence results for large families of Bayesian games is clearly an important future direction of research. Since, as we argued before, most games are Bayesian games, the existence of a Bayesian equilibrium should, and could, reach at least the level of generality available for the existence of a Nash equilibrium.

Acknowledgments I am grateful to two anonymous reviewers for their helpful comments.

Bibliography

- Aumann R (1974) Subjectivity and correlation in randomized strategies. *J Math Econ* 1:67–96
- Aumann R (1976) Agreeing to disagree. *Ann Stat* 4:1236–1239
- Aumann R (1987) Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55:1–18
- Aumann R (1998) Common priors: a reply to Gul. *Econometrica* 66:929–938
- Aumann R (1999a) Interactive epistemology I: knowledge. *Int J Game Theory* 28:263–300
- Aumann R (1999b) Interactive epistemology II: probability. *Int J Game Theory* 28:301–314
- Aumann R, Heifetz A (2002) Incomplete information. In: Aumann R, Hart S (eds) *Handbook of game theory*, vol 3. Elsevier, North Holland, pp 1666–1686
- Aumann R, Maschler M (1995) *Repeated games with incomplete information*. MIT Press, Cambridge
- Brandenburger A, Dekel E (1993) Hierarchies of beliefs and common knowledge. *J Econ Theory* 59:189–198
- Gul F (1998) A comment on Aumann's Bayesian view. *Econometrica* 66:923–927
- Harsanyi J (1967–1968) Games with incomplete information played by 'Bayesian' players, parts I–III. *Manag Sci* 8:159–182, 320–334, 486–502
- Heifetz A (1993) The Bayesian formulation of incomplete information, the non-compact case. *Int J Game Theory* 21:329–338
- Heifetz A, Mongin P (2001) Probability logic for type spaces. *Games Econ Behav* 35:31–53
- Heifetz A, Samet D (1998) Topology-free topology of beliefs. *J Econ Theory* 82:324–341
- Maskin E, Riley J (2000) Asymmetric auctions. *Rev Econ Stud* 67:413–438
- Meier M (2001) An infinitary probability logic for type spaces. CORE discussion paper 2001/61
- Mertens J-F, Sorin S, Zamir S (1994) Repeated games, part A: background material. CORE discussion paper no 9420
- Mertens J-F, Zamir S (1985) Foundation of Bayesian analysis for games with incomplete information. *Int J Game Theory* 14:1–29
- Milgrom PR, Stokey N (1982) Information, trade and common knowledge. *J Eco Theory* 26:17–27
- Milgrom PR, Weber RJ (1982) A theory of auctions and competitive bidding. *Econometrica* 50:1089–1122
- Nyarko Y (1991) Most games violate the Harsanyi doctrine. C.V. Starr working paper #91–39, NYU
- Reny P, Zamir S (2004) On the existence of pure strategy monotone equilibria in asymmetric first price auctions. *Econometrica* 72:1105–1125
- Sorin S, Zamir S (1985) A 2-person game with lack of information on $1\frac{1}{2}$ sides. *Math Oper Res* 10:17–23
- Vassilakis S, Zamir S (1993) Common beliefs and common knowledge. *J Math Econ* 22:495–505
- Wolfstetter E (1999) *Topics in microeconomics*. Cambridge University Press, Cambridge

Repeated Games with Complete Information

Olivier Gossner¹ and Tristan Tomala²

¹PSE, UMR CNRS-EHESS-ENPC-ENS 8545, Northwestern University, Paris, France

²Economics and Finance Department, HEC Paris, Paris, France

Article Outline

Glossary

Definition of the Subject

Introduction

Games with Observable Actions

Games with Non-observable Actions

Bibliography

Glossary

Behavioral strategy A decision rule that prescribes a randomized choice of actions for every possible history.

Monitoring structure A description of player's observation of each other's strategic choices. It specifies, for every profile of actions, the probability distribution over the profiles of individual signals received by the agents.

Nash equilibrium A strategy profile from which no unilateral deviation is profitable.

Repeated game A model of repeated interaction between agents.

Sequential equilibrium A strategy profile and Bayesian beliefs on past histories such that after every history, every agent is acting optimally given his beliefs.

Definition of the Subject

Repeated interactions arise in several domains such as Economics, Computer Science, and Biology.

The theory of repeated games models situations in which a group of agents engage in a strategic interaction over and over. The data of the strategic interaction is fixed over time and is known by all the players. This is in contrast with stochastic games, for which the data of the strategic interaction is controlled by player's choices, and repeated games with incomplete information, where the stage game is not common knowledge among players (the reader is referred to the corresponding entries of this encyclopedia). Early studies of repeated games include Luce and Raiffa (1957) and Aumann (1960). In the context of production games, Friedman (1971) shows that, while the competitive outcome is the only one compatible with individual profit maximization under a static interaction, collusion is sustainable at an equilibrium when the interaction is repeated.

Generally, repeated games provide a framework in which individual utility maximization by selfish agents is compatible with welfare maximization (common good), while this is known to fail for many classes of static interactions.

Introduction

The discussion of an example shows the importance of repeated games and introduces the questions studied.

Consider the following game referred to as the Prisoner's Dilemma:

	C	D	
C	4, 4	5, 5	The Prisoner's Dilemma
D	5, 0	1, 1	

Player 1 chooses the row, player 2 chooses the column, and the pair of numbers in the corresponding cell are the payoffs to players 1 and 2 respectively.

In a one-shot interaction, the only outcome consistent with game theory predictions is (D, D). In fact, each player is better off playing D whatever the other player does.

On the other hand, if players engage in a repeated Prisoner's Dilemma, if they value sufficiently future payoffs compared to present ones, and if past actions are observable, then (C, C) is a sustainable outcome. Indeed, if each player plays C as long as the other one has always done so in the past and plays D otherwise, both players have an incentive to always play C , since the short-term gain that can be obtained by playing D is more than offset by the future losses entailed by the opponent playing D at all future stages.

Hence, a game theoretical analysis predicts significantly different outcomes from a repeated game than from static interaction. In particular, in the Prisoner's Dilemma, the cooperative outcome (C, C) can be sustained in the repeated game, while only the noncooperative outcome (D, D) can be sustained in one-shot interactions.

In general, what are the equilibrium payoffs of a repeated game and how can they be computed from the data of the static game? Is there a significant difference between games repeated a finite number of times and infinitely repeated ones? What is the role played by the degree of impatience of players? Do the conclusions obtained for the Prisoner's Dilemma game and for other games rely crucially on the assumption that each player perfectly observes other player's past choices, or would imperfect observation be sufficient? The theory of repeated games aims at answering these questions and many more.

Games with Observable Actions

This section focuses on repeated games with perfect monitoring in which, after every period of the repeated game, all strategic choices of all the players are publicly revealed.

Data of the Game, Strategies, and Payoffs

Data of the Stage Game

There is a finite set I of players. A stage game is repeated over and over. Each player i 's action set in this stage game is denoted A_i , and $S_i = \Delta(A_i)$ is the set of player i 's mixed actions (for any finite set X , $\Delta(X)$ denotes the set of probabilities over X).

Every degenerate lottery in S_i (which puts probability 1 to one particular action in A_i) is associated to the corresponding element in A_i . A choice of action for every player i determines an outcome $a \in \prod_i A_i$. The payoff function of the stage game is

$g : A \rightarrow \mathbb{R}^I$. Payoffs are naturally associated to profiles of mixed actions $s \in S = \prod_i S_i$ using the expectation: $g(s) = \mathbf{E}_s g(a)$.

Repeated Game

After every repetition of the stage game, the action profile previously chosen by the players is publicly revealed. After the t first repetitions of the game, a player's information consists of the publicly known history at stage t , which is an element of $H_t = A^t$ ($H_0 = \{\emptyset\}$ by convention). A strategy in the repeated game specifies the choice of a mixed action at every stage, as a function of the past observed history. More specifically, a *behavioral* strategy for player i is of the form $\sigma_i : \cup_t H_t \rightarrow S_i$. When all the strategy choices belong to A_i ($\sigma_i : \cup_t H_t \rightarrow A_i$), σ_i is called a *pure* strategy.

Other Strategy Specifications A behavioral strategy allows the player to randomize his action depending on past history. If, at the start of the repeated game, the player was to randomize over the set of behavioral strategies, the result would be equivalent to a particular behavioral strategy choice. This result is a consequence of Kuhn's theorem (Aumann 1964; Kuhn 1953). Furthermore, behavioral strategies are also equivalent to randomizations over the set of pure strategies.

Induced Plays Every choice of pure strategies $\sigma = (\sigma_i)_i$ by all the players induces a play $h = (a_1, a_2, \dots) \in A^\infty$ in the repeated game, defined inductively by $a_1 = (\sigma_{i,0}(\emptyset))$ and $a_t = (\sigma_{i,t-1}(a_1, \dots, a_{t-1}))$. A profile of behavioral strategies σ defines a probability distribution P_σ over plays.

Preferences To complete the definition of the repeated game, it remains to define player's preferences over plays. The literature commonly distinguishes infinitely repeated games with or without discounting and finitely repeated games.

In infinitely repeated games with no discounting, the players care about their long-run stream of stage payoffs. In particular, the payoff in the repeated game associated to a play $h = (a_1, a_2, \dots) \in A^\infty$ coincides with the limit of the Cesaro means of stage payoffs when this limit exists. When this limit does not exist, the most common evaluation of the stream of payoffs is defined through a Banach limit of the Cesaro means (a Banach limit is a linear form on the set of bounded sequences that lies always between the liminf and the limsup).

In infinitely repeated games with discounting, a discount factor $0 < \delta < 1$ characterizes the player's degree of impatience. A payoff of 1 at stage $t + 1$ is equivalent to a payoff of δ at stage t . Player i 's payoff in the repeated game for the play $h = (a_1, a_2, \dots) \in A^\infty$ is the normalized sum of discounted payoffs: $(1 - \delta) \sum_{t \geq 1} \delta^{t-1} g_i(a_t)$.

In finitely repeated games, the game ends after some stage T . Payoffs induced by the play after stage T are irrelevant (and a strategy needs not specify choices after stage T). The payoff for a player is the average of the stage payoffs during stages 1 up to T : $\frac{1}{T} \sum_{t=1}^T g_i(a_t)$.

Equilibrium Notions What plays can be expected to be observed in repeated interactions of players who observe each other's choices? Noncooperative game theory focuses mainly on the idea of stable convention, i.e., of strategy profiles from which no player has incentives to deviate, knowing the strategies adopted by the other players.

A strategy profile forms a Nash equilibrium (Nash 1951) when no player can improve his payoff by choosing an alternative strategy, as long as other players follow the prescribed strategies.

In some cases, the observation of past play may not be consistent with the prescribed strategies. When, for every possible history, each player's strategy maximizes the continuation stream of payoffs, assuming that other players abide with their prescribed strategies at all future stages, the strategy profile forms a subgame perfect equilibrium (Selten 1965).

Perfect equilibrium is a more robust and often considered a more satisfactory solution concept than Nash equilibrium. The construction of perfect equilibria is in general also more demanding than the construction of Nash equilibria.

The main objective of the theory of repeated games is to characterize the set of payoff vectors that can be sustained by some Nash or perfect equilibrium of the repeated game.

Necessary Conditions on Equilibrium Payoffs

Some properties are common to all equilibrium payoffs. First, under the common assumption that all players evaluate the payoff associated to a play in the same way, the resulting payoff vector in the repeated game is a convex combination of stage payoffs. That is, the payoff vector in the repeated game is an element of the convex closure of $g(A)$, called the set of *feasible payoffs* and denoted F .

A notable exception is the work of Lehrer and Pauzner (1999) who study repeated games where players have heterogeneous time preferences. The payoff vector resulting from a play does not necessarily belong to F if players have different evaluations of payoff streams. For instance, in a repetition of the Prisoner's Dilemma, if player 1 cares only about the payoff in stage 1 and player 2 cares only about the payoff in stage 2, it is possible for both players to obtain a payoff of four in the repeated game.

Now consider a strategy profile σ , and let τ_i be a strategy of player i that plays after every history $(a_1, \dots, a_t)$ a best response to the profile of mixed actions chosen by the other players in the next stage. At any stage of the repeated game, the expected payoff for player i using τ_i is no less than

$$v_i = \min_{s_{-i} \in S_{-i}} \max_{a_i \in A_i} g_i(s_{-i}, a_i) \quad (1)$$

where $s_{-i} = (s_j)_{j \neq i}$ (we use similar notations throughout the paper: for a family of sets $(E_i)_{i \in I}$, e_{-i} denotes an element of $E_{-i} =$

$\prod_{j \neq i} E_j$, and a profile $e \in \prod_j E_j$ is denoted $e = (e_i, e_{-i})$ when the i th component is stressed).

The payoff v_i is referred to as player i 's min max payoff. A payoff vector that provides each player i with at least [resp. strictly more than] v_i is called *individually rational* [resp. strictly individually rational], and IR [resp. IR^*] denotes the set of such payoff vectors. Since, for any strategy profile, there exists a strategy of player i that yields a payoff no less than v_i , all equilibrium payoffs have to be individually rational.


Also note that players $j \neq i$ collectively have a strategy profile in the repeated game that forces player i 's payoff down to v_i : they play repeatedly a mixed strategy profile that achieves the minimum in the definition of v_i . Such a strategy profile in the one-shot game is referred to as punishing strategy or min max strategy against player i .

For the Prisoner's Dilemma game, F is the convex hull of $(1, 1)$, $(5, 0)$, $(0, 5)$, and $(4, 4)$. Both player's min max levels are equal to 1. Figure 1 illustrates the set of feasible and individually rational payoff vectors (hatched area):

The set of feasible and individually rational payoffs can be directly computed from the stage game data.

Infinitely Patient Players

The following result has been part of the folklore of game theory at least since the mid-1960s. Its authorship is obscure (see the introduction of Aumann (1981a)). For this reason, it is commonly referred to as the "Folk Theorem." By extension, characterization of sets of equilibrium payoffs in repeated games is also referred to as "Folk Theorems."



Repeated Games with Complete Information,
Fig. 1 F and IR for the Prisoner's Dilemma

Theorem 1 The set of equilibrium payoffs of the repeated game with no discounting coincides with the set of feasible and individually rational payoffs.

Aumann and Shapley (1976, 1994) and Rubinstein (1977, 1994) show that restricting attention to perfect equilibria does not narrow down the set of equilibrium payoffs. They prove that

Theorem 2 The set of perfect equilibrium payoffs of the repeated game with no discounting coincides with the set of feasible and individually rational payoffs.

We outline a proof of Theorem 2. It is established that any equilibrium payoff is in $F \cap IR$. We need only to prove that every element of $F \cap IR$ is a subgame perfect equilibrium payoff. Let $x \in F \cap IR$, and let $h = a_1, \dots, a_t, \dots$ be a play inducing x . Consider the strategies that play a_t in stage t ; if player i does not respect this prescription at stage t_0 , the other players punish player i for t_0 stages by repeatedly playing the min max strategy profile against player i . After the punishment phase is over, players revert to the play of h , hence playing $a_{2t_0+1} \dots$

Now we explain why these strategies form a subgame perfect equilibrium. Consider a strategy of player i starting after any history. The induced play by this strategy for player i and by other player's prescribed strategies is, up to a subset of stages of null density, defined by the sequence h with interweaved periods of punishment for player i . Hence the induced long-run payoff for player i is a convex combination of the punishment payoff and of the payoff induced by h . The result follows since the payoff for player i induced by h is no worse than the punishment payoff.

Impatient Players

The strategies constructed in the proof of the Folk Theorem for repeated games with infinitely patient players (Theorem 1) do not necessarily constitute a subgame perfect equilibrium if players are impatient. Indeed, during a punishment phase, the punishing players may be

receiving low stage payoffs, and these stage payoffs matter in the evaluation of their stream of payoffs. When constructing subgame perfect equilibria of discounted games, one must make sure that after a deviation of player i , players $j \neq i$ have incentives to implement player i 's punishment.

Nash Reversion

Friedman (1971) shows that every feasible payoff that Pareto dominates a Nash equilibrium payoff of the static game is a subgame perfect equilibrium payoff of the repeated game provided that players are patient enough. In Friedman's proof, punishments take the simple form of reversion to the repeated play of the static Nash equilibrium forever. In the Prisoner's Dilemma, (D, D) is the only static Nash equilibrium payoff, and thus $(4, 4)$ is a subgame perfect Nash equilibrium payoff of the repeated game if players are patient enough. Note however that in some games, the set of payoffs that Pareto dominates some equilibrium payoff may be empty. Also, Friedman's result constitutes a partial Folk Theorem only in that it does not characterize the full set of equilibrium payoffs.

The Recursive Structure

Repeated games with discounting possess a structure similar to dynamic programming problems. At any stage in time, players choose actions that maximize the sum of the current payoff and the payoff at the subsequent stages. When strategies form a subgame perfect equilibrium, the payoff vector at subsequent stages must be an equilibrium payoff, and players must have incentives to follow the prescribed strategies at the current stage. This implies that subgame perfect equilibrium payoffs have a recursive structure, first studied by Abreu (1988). Subsection "A Recursive Structure" presents the recursive structure in more details for the more general model of games with public monitoring.

The Folk Theorem for Discounted Games

Relying on Abreu's recursive results, Fudenberg and Maskin (1986) prove the following Folk Theorem for subgame perfect equilibria with discounting:

Theorem 3 If the number of players is 2 or if the set feasible payoff vectors has a nonempty interior, then any payoff vector that is feasible and strictly individually rational is a subgame perfect equilibrium of the discounted repeated game, provided that players are sufficiently patient.

Forges et al. (1986) provide an example for which a payoff which is individually rational but not strictly individually rational is not an equilibrium payoff of the discounted game.

Abreu et al. (1994) show that the nonempty interior condition of the theorem can be replaced by a weaker condition of "nonequivalent utilities": no pair of players have the same preferences over outcomes. Wen (1994) and Fudenberg et al. (2007) show that a Folk Theorem still holds when the condition of nonequivalent utilities fails if one replaces the min max level defining individually rational payoffs by some "effective min max" payoffs.

An alternative representation of impatience to discounted payoffs in infinitely repeated games is the overtaking criterion, introduced by Rubinstein (1979): the play $(a_1, a_2, \dots)$ is strictly preferred by player i to the play $(a'_1, a'_2, \dots)$ if the inferior limit of the difference of the corresponding streams of payoffs is positive, i.e., if

$$\liminf_T \sum_{t=1}^T g_i(a_t) - g_i(a'_t) > 0 \quad \text{Rubinstein}$$

(1979) proves a Folk Theorem with the overtaking criterion.

Finitely Repeated Games

Strikingly, equilibrium payoffs in finitely repeated games and in infinitely repeated games can be drastically different. This effect is best exemplified in repetitions of the Prisoner's Dilemma.

The Prisoner's Dilemma

Recall that in an infinitely repeated Prisoner's Dilemma, cooperation at all stages is achieved at a subgame perfect equilibrium if players are patient enough. By contrast, at every Nash equilibrium of any finite repetition of the Prisoner's Dilemma, both players play D at every stage with probability 1.

Now we present a short proof of this result. Consider any Nash equilibrium of the Prisoner's Dilemma repeated T times. Let $a_1, \dots, a_T$ be a sequence of action profiles played with positive probability at the Nash equilibrium. Since each player can play D at the last stage of the repetition, and D is a dominating action, $a_T = (D, D)$. We now prove by induction on τ that for any such τ , $a_{T-\tau}, \dots, a_T = (D, D), \dots, (D, D)$. Assume the induction hypothesis valid for $\tau - 1$. Consider a strategy for player i that follows the equilibrium strategy up to stage $T - \tau - 1$, then plays D from stage $T - \tau$ on. This strategy obtains the same payoff as the equilibrium strategy at stages $1, \dots, T - \tau - 1$, and at least as much as the equilibrium strategy at stages $T - \tau + 1, \dots, T - \tau$. Hence, this strategy cannot obtain more than the equilibrium strategy at stage $T - \tau$, and therefore, the equilibrium strategy plays D at stage $T - \tau$ with probability 1 as well.

Sorin (1986) proves the more general result:

Theorem 4 Assume that in every Nash equilibrium of G , all players are receiving their individually rational levels. Then, at every Nash equilibrium of any finitely repeated version of G , all players are receiving their individually rational levels.

The proof of Theorem 4 relies on a backward induction type of argument, but it is striking that the result applies for *all* Nash equilibria and not only for *subgame perfect* Nash equilibria. This result shows that, unless some additional assumptions are made on the one-shot game, a Folk Theorem cannot obtain for finitely repeated games.

Games with Unique Nash Payoff

Using a proof by backward induction, Benoît and Krishna (1985) obtain the following result:

Theorem 5 Assume that G admits x as unique Nash equilibrium payoff. Then every finite repetition of G admits x as unique subgame perfect equilibrium payoff.

Theorems 4 and 5 rely on the assumption that the last stage of repetition, T , is common

knowledge between players. Neyman (1999) shows that a Folk Theorem obtains for the finitely repeated Prisoner's Dilemma (and for other games) if there is lack of common knowledge on the last stage of repetition.

Folk Theorems for Finitely Repeated Games

A Folk Theorem can be obtained when there are two Nash equilibrium payoffs for each player. The following result is due to Benoît and Krishna (1985) and Gossner (1995):

Theorem 6 Assume that each player has two distinct Nash equilibrium payoffs in G and that the set of feasible payoffs has nonempty interior. Then, the set of subgame perfect equilibrium payoffs of the T times repetition of G converges to the set of feasible and individually rational payoffs as T goes to infinity.

Hence, with at least two equilibrium payoffs per player, the sets of equilibrium payoffs of finitely repeated games and infinitely repeated games are asymptotically the same.

The condition that each player has two distinct Nash equilibrium payoffs in the stage game can be weakened; see Smith (1995). Assume for simplicity that *one* player has two distinct Nash payoffs. By playing one of the two Nash equilibria in the last stages of the repeated game, it is possible to provide incentives for this player to play actions that are not part of Nash equilibria of the one-shot game in previous stages. If this construction leads to perfect equilibria in which a player $j \neq i$ has distinct payoffs, we can now provide incentives for both players i and j . If successive iterations of this procedure yield distinct subgame perfect equilibrium payoffs for all players, a Folk Theorem applies.

Games with Non-observable Actions

For infinitely repeated games with perfect monitoring, a complete and simple characterization of the set of equilibrium payoffs is obtained: feasible and individually rational payoff vectors. In particular, cooperation can be sustained at equilibrium.

How equilibrium payoffs of the repeated game depend on the quality of player's monitoring of each other's actions is the subject of a very active area of research.

Repeated games with imperfect monitoring, in which players observe imperfectly other player's action choices, were first motivated by economic applications. In Stigler (1964), two firms are repeatedly engaged in price competition over market shares. Each firm observes its own sales, but not the price set by the rival. While it is in the best interest for both firms to set a collusive price, each firm has incentives to secretly undercut the rival's price. Upon observing plunging sales, should a firm deduce that the rival firm is undercutting prices, and retaliate by setting lower prices, or should lower sales be interpreted as a result of an exogenous shock on market demand? Whether collusive behavior is sustainable or not at equilibrium is one of the motivating questions in the theory of repeated games with imperfect monitoring.

It is interesting to compare repeated games with imperfect monitoring with their perfect monitoring counterparts.

The structure of equilibria used to prove the Folk Theorem with perfect monitoring and no discounting is rather simple: if a player deviates from the prescribed strategies, the deviation is detected, and the deviating player is identified, and all other players can then punish the deviator. With imperfect monitoring, not all deviations are detectable, and when a deviation is detected, deviators are not necessarily identifiable. The notions of detection and identification allow fairly general Folk Theorems for undiscounted games. We present these results in subsection “[Detection and Identification](#).”

With discounting, repeated games with perfect monitoring possess a recursive structure that facilitates their study. Recursive methods can also be successfully applied to discounted games with public monitoring. We review the major results of this branch of the literature in subsection “[Public Equilibria](#).”

Almost-perfect monitoring is the natural framework to study the effect of small departures from the perfect or public monitoring

assumptions. We review this literature in subsection “[Almost-Perfect Monitoring](#).”

Little is known about general discounted games with imperfect private monitoring. We present the main known results in subsection “[General Stochastic Signals](#).”

With perfect monitoring, the worst equilibrium payoff for a player is given by the min max of the one-shot game, where punishing (minimizing) players choose an independent profile of mixed strategies. With imperfect monitoring, correlation past signals for the punishing players may lead to more efficient punishments. We present results on punishment levels in subsection “[Punishment Levels](#).”

Model

In this section we define repeated games with imperfect monitoring and describe several classes of monitoring structures of particular interest.

Data of the Game

Recall that the one-shot strategic interaction is described by a finite set I of players, a finite action set A_i for each player i , and a payoff function $g : A \rightarrow \mathbb{R}^I$. Player's observation of each other's actions is described by a *monitoring structure* given by a finite set of signals Y_i for each player i and by a transition probability $Q : A \rightarrow \Delta(Y)$ (with $A = \prod_{i \in I} A_i$ and $Y = \prod_{i \in I} Y_i$). When the action profile chosen is $a = (a_i)_{i \in I}$, a profile of signals $y = (y_i)_{i \in I}$ is drawn with probability $Q(y|a)$, and y_i is observed by player i .

Perfect Monitoring Perfect monitoring is the particular case in which each player observes the action profile chosen: for each player i , $Y_i = A$ and $Q((y_i)_{i \in I} | a) = 1_{\{\forall i, y_i = a\}}$.

Almost-Perfect Monitoring The monitoring structure is ε -perfect (see Mailath and Morris (2002)) when each player can identify the other player's action with a probability of error less than ε . This is the case if there exist functions $f_i : A_i \times Y_i \rightarrow A_-i$ for all i such that for all $a \in A$:

$$Q(\forall_i, f_i(a_i, y_i) = a_{-i}|a) \geq 1 - \varepsilon.$$

Almost-perfect monitoring refers to ε -perfect monitoring for small values of ε .

Canonical Structure The monitoring structure is *canonical* when each player's observation corresponds to an action profile of the opponents, that is, when $Y_i = A_{-i}$.

Public and Almost-Public Signals Signals are *public* when all the players observe the same signal, i.e., $Q(\forall_i, j, y_i = y_j|a) = 1$, for every action profile a . For instance, in Green and Porter (1984), firms compete over quantities, and the public signal is the realization of the price. Firms can then make inferences on rival's quantities based on their own quantity and market price.

The case in which $Q(\forall_i, j, y_i = y_j|a)$ is close to 1 for every a is referred to as *almost-public monitoring* (see Mailath and Morris (2002)).

Private signals refer to the case when these signals are not public.

Deterministic Signals Signals are *deterministic* when the signal profile is uniquely determined by the action profile. When a is played, the signal profile y is given by $y = f(a)$, where f is called the *signaling function*.

Observable Payoffs Payoffs are *observable* when each player i can deduce his payoff from his action and his signal. This is the case if there exists a mapping $\varphi : A_i \times Y_i \rightarrow \mathbb{R}$ such that for every action profile a , $Q(\forall_i, g_i(a) = \varphi(a_i, y_i|a) = 1$.

The Repeated Game

The game is played repeatedly, and after every stage t , the profile of signals y_t received by the players is drawn according to the distribution $Q(y_t|a_t)$, where a_t is the profile of action chosen at stage t . A player's information consists of his past actions and signals. We let $H_{i,t} = (A_i \times Y_i)^t$ be the set of *player i 's histories* of length t . A strategy for player i now consists of a mapping $\sigma_i : \cup_{t \geq 0} H_{i,t} \rightarrow S_i$. The set of complete histories of

the game after t stages is $H_t = (A \times Y)^t$; it describes chosen actions and received signals for all the players at all past stages. A strategy profile $\sigma = (\sigma_i)_{i \in I}$ induces a probability distribution P_σ on the set of plays $H_\infty = (A \times Y)^\infty$.

Equilibrium Notions

Nash Equilibria Player's preferences over game plays are defined according to the same criteria as for perfect monitoring. We focus on infinitely repeated games, both discounted and undiscounted. A choice of players' preferences defines a set of Nash equilibrium payoffs in the repeated game.

Sequential Equilibria The most commonly used refinement of Nash equilibrium for repeated games with imperfect monitoring is the *sequential equilibrium* concept (Kreps and Wilson 1982), which we recall here.

A *belief assessment* is a sequence $\mu = (\mu_{i,t})_{t \geq 1, i \in I}$ with $\mu_{i,t} : H_{i,t} \rightarrow \Delta(H_t)$, i.e., given the private history h_i of player i , $\mu_{i,t}(h_i)$ is the probability distribution representing the belief that player i holds on the full history.

A sequential equilibrium of the repeated game is a pair (σ, μ) where σ is a strategy profile and μ is a belief assessment such that (1) for each player i and every history h_i , σ_i is a best reply in the continuation game, given the strategies of the other players and the belief that player i holds regarding the past, and (2) the beliefs must be consistent in the sense that (σ, μ) is the limit of a sequence (σ^n, μ^n) where for every n , σ^n is a completely mixed strategy (it assigns positive probability to every action after every history) and μ^n is the unique belief derived from Bayes' law under P_{σ^n} .

Sequential equilibria are defined both on the discounted game and the undiscounted versions of the repeated game.

For undiscounted games, the set of Nash equilibrium payoffs and sequential equilibrium payoffs coincides. The two notions also coincide for discounted games when the monitoring has full

support (i.e., under every action profile, all signal profiles have positive probability). The results presented in this survey all hold for sequential equilibria, both for discounted and undiscounted games.

Extensions of the Repeated Game

When players receive correlated inputs or may communicate between stages of the repeated game, the relevant concepts are correlated and communication equilibria.

Correlated Equilibria A correlated equilibrium (Aumann 1974) of the repeated game is an equilibrium of an extended game in which at a preliminary stage, a mediator chooses a profile of correlated random inputs and informs each player of his own input; then the repeated game is played. A characterization of the set of correlated equilibrium payoffs for two-player games is obtained by Lehrer (1992a).

Correlation arises endogenously in repeated games with imperfect monitoring, as the signals received by the players can serve as correlated inputs that influence player's continuation strategies. This phenomenon is called *internal correlation* and was studied by Lehrer (1991) and Gossner and Tomala (2006, 2007).

Communication Equilibria An (extensive form) communication equilibrium (Myerson 1982; Forges 1986) of a repeated game is an equilibrium of an extension of the repeated game in which after every stage, players send messages to a mediator and the mediator sends back private outputs to the players. Characterizations of the set of communication equilibrium payoffs are obtained under weak conditions on the monitoring structure; see, e.g., Kandori and Matsushima (1998), Compte (1998), and Renault and Tomala (2004).

Detection and Identification

Equivalent Actions

A player's deviation is detectable when it induces a different distribution of signals for other players. When two actions induce the same distribution of

actions for other players, they are called equivalent (Lehrer 1990, 1991, 1992a, b):

Definition 1 Two actions a_i and b_i of player i are equivalent, and we note $a_i \sim b_i$, if they induce the same distribution of other players' signals:

$$Q(y_{-i}|a_i, a_{-i}) = Q(y_{-i}|b_i, a_{-i}), \quad \forall a_{-i}.$$

Example 1 Consider the two-player repeated Prisoner's Dilemma where player 2 receives no information about the actions of player 1 (e.g., Y_2 is a singleton). The two actions of player 1 are thus equivalent. The actions of player 2 are independent of the actions of player 1: player 1 has no impact on the behavior of player 2. Player 2 has no power to threat player 1 and in any equilibrium, player 1 defects at every stage. Player 2 also defects at every stage: since player 1 always defects, he also loses his threatening power. The only equilibrium payoff in this repeated game is thus (1, 1).

Example 1 suggests that between two equivalent actions, a player chooses at equilibrium the one that yields the highest stage payoff. This is indeed the case when the information received by a player does not depend on his own action. Lehrer (1990) studies particular monitoring structures satisfying this requirement. Recall from Lehrer (1990) the definition of *semi-standard* monitoring structures: each action set A_i is endowed with a partition $\bar{A}_i$; when player i plays a_i , the corresponding partition cell $\bar{a}_i$ is publicly announced. In the semi-standard case, two actions are equivalent if and only if they belong to the same cell: $a_i \sim b_i \Leftrightarrow \bar{a}_i = \bar{b}_i$ and the information received by a player on other player's action does not depend on his own action.

If player i deviates from a_i to b_i , the deviation is undetected if and only if $a_i \sim b_i$. Otherwise it is detected by all other players. A profile of mixed actions is called *immune to undetectable deviations* if no player can profit by a unilateral deviation that maintains the same distribution of other players' signals. The following result, due to Lehrer (1990), characterizes equilibrium payoffs for undiscounted games with semi-standard signals:

Theorem 7 In an undiscounted repeated game with semi-standard signals, the equilibrium payoffs are the individually rational payoffs that belong to the convex hull of payoffs generated by mixed action profiles that are immune to undetectable deviations.

More Informative Actions

When the information of player i depends on his own action, some deviations may be detected in the course of the repeated game even though they are undetectable in the stage game.

Example 2 Consider the following modification of the Prisoner's Dilemma. The action set of player 1 is $A_1 = \{C_1, D_1\} \times \{C_2, D_2\}$ and the action set of player 2 is $\{C_2, D_2\}$. An action for player 1 is thus a pair $a_1 = (\tilde{a}_1, \tilde{a}_2)$. When the action profile $(\tilde{a}_1, \tilde{a}_2, a_2)$ is played, the payoff to player i is $g_i(\tilde{a}_1, a_2)$. We can interpret the component $\tilde{a}_1$ as a *real* action (it impacts payoffs) and the component $\tilde{a}_2$ as a *message* sent to player 2 (it does not impact payoffs). The monitoring structure is as follows:

- Player 2 only observes the message component $\tilde{a}_2$ of the action of player 1.
- Player 1 perfectly observes the action of player 2 if he chooses the cooperative real action ($\tilde{a}_1 = C_1$) and gets no information on player 2's action if he defects ($\tilde{a}_1 = D_1$).

Note that the actions (C_1, C_2) and (D_1, C_2) of player 1 are equivalent, and so are the actions (C_1, D_2) and (D_1, D_2) . However, it is possible to construct an equilibrium that implements the cooperative payoff along the following lines:

1. Using his message component, player 1 reports at every stage $t > 1$ the previous action of player 2. Player 1 is punished in case of a nonmatching report.
2. Player 2 randomizes between both actions, so that player 1 needs to play the cooperative action in order to report player 2's action accurately. The weight on the defective action of player 2 goes to 0 as t goes to infinity to ensure efficiency.

Player 2 has incentives to play C_2 most of the time, since player 1 can statistically detect if player 2 uses the action D_2 more frequently than prescribed. Player 1 also has incentives to play the real action C_1 , as this is the only way to observe player 2's action, which needs to be reported later on.

The key point in the example above is that the two real actions C_1 and D_1 of player 1 are equivalent but D_1 is *less informative* than C_1 for player 1. For general monitoring structures an action a_i is *more informative* than an action b_i if, *whenever player i plays a_i i , he can reconstitute the signal he would have observed, had he played b_i* . The precise definition of the *more informative* relation relies on Blackwell's ordering of stochastic experiments (Blackwell 1951):

Definition 2 The action a_i of player i is more informative than the action b_i if there exists a transition probability $f : Y_i \rightarrow \Delta(Y_i)$ such that for every a_{-i} and every profile of signals y ,

$$\sum_{y_i} f(y'_i | y_i) Q(y_i, y_{-i} | a_i, a_{-i}) = Q(y'_i, y_{-i} | b_i, a_{-i}).$$

We denote $a_i \succ b_i$ if $a_i \sim b_i$ and a_i is more informative than b_i .

Assume that prescribed strategies require player i to play b_i at stage t , and let $a_i \succ b_i$. Consider the following deviation from player i : play a_i at stage t , and reconstruct a signal at stage t that could have arisen from the play of b_i . In all subsequent stages, play as if no deviation took place at stage t and as if the reconstructed signal had been observed at stage t . Not only such a deviation would be undetectable at stage t , since $a_i \sim b_i$, but it would also be undetectable at all subsequent stages, as it would induce the same probability distribution over plays as under the prescribed strategy. This argument shows that, if an equilibrium strategy specifies that player i plays a_i , there is no $b_i \succ a_i$ that yields a higher expected stage payoff than a_i .

Definition 3 A distribution of action profiles $p \in \Delta(A)$ is *immune to undetectable deviations* if for each player i and pair of actions a_i, b_i such that $b_i \succ a_i$:

$$\sum_{a_{-i}} p(a_i, a_{-i}) g_i(a_i, a_{-i}) \geq \sum_{a_{-i}} p(a_i, a_{-i}) g_i(b_i, a_{-i})$$

If p is immune to undetectable deviations and if player i is supposed to play a_i , any alternative action b_i that yields a greater expected payoff cannot be such that $b_i \succ a_i$.

The following proposition gives a necessary condition on equilibrium payoffs that holds both in the discounted and in the undiscounted cases:

Proposition 1 Every equilibrium payoff of the repeated game is induced by a distribution that is immune to undetectable deviations.

The condition of Proposition 1 is tight for some specific classes of games, all of them assuming two players and no discounting.

Following Lehrer (1992a), signals are *nontrivial* if, for each player i , there exist an action a_i for player i and two actions a_j and b_j for i 's opponent such that the signal for player i is different under (a_i, a_j) and (a_i, b_j) . Lehrer (1992a) proves:

Theorem 8 The set of correlated equilibrium payoffs of the undiscounted game with deterministic and nontrivial signals is the set of individually rational payoffs induced by distributions that are immune to undetectable deviations.

Lehrer (1992b) assumes that payoffs are observable and obtains the following result:

Theorem 9 In a two-player repeated game with no discounting, nontrivial signals, and observable payoffs, the set of equilibrium payoffs is the set of individually rational payoffs induced by distributions that are immune to undetectable deviations.

Finally, Lehrer (1991) shows that, in some cases, one may dispense with the correlation device of Theorem 8, as all necessary correlation can be generated endogenously through the signals of the repeated game:

Proposition 2 In two-player games with nontrivial signals such that either the action profile is publicly announced or a blank signal is publicly announced, the set of equilibrium payoffs coincides with the set of correlated equilibrium payoffs.

Identification of Deviators

A deviation is identifiable when every player can infer the identity of the deviating player from his observations. For instance, in a game with public signals, if separate deviations from players i and j induce the same distribution of public signals, these deviations from i or j are not identifiable. In order to be able to punish the deviating player, it is sometimes necessary to know his identity. Detectability and identifiability are two separate issues, as shown by the following example:

Example 3 Consider the following three-player game where player 1 chooses the row, player 2 chooses the column, and player 3 chooses the matrix.

		<i>L</i>	<i>R</i>		<i>L</i>	<i>R</i>		<i>L</i>	<i>R</i>
<i>T</i>		1, 1, 1	4, 4, 0		0, 3, 0	0, 3, 0		3, 0, 0	3, 0, 0
<i>B</i>		4, 4, 0	4, 4, 0		0, 3, 0	0, 3, 0		3, 0, 0	3, 0, 0
		<i>W</i>			<i>M</i>			<i>E</i>	

Consider the monitoring structure in which actions are not observable and the payoff vector is publicly announced.

The payoff $(1, 1, 1)$ is feasible and individually rational. The associated action profile (T, L, W) is immune to undetectable deviations since any individual deviation from (T, L, W) changes the payoff.

However, $(1, 1, 1)$ is not an equilibrium payoff. The reason is that, player 3, who has the power to punish either player 1 or player 2, cannot punish both players simultaneously: punishing player 1 rewards player 2 and vice versa. More precisely, whatever weights player 3 puts on the action M and E , the sum of player 1 and player 2's payoffs is greater than 3. Any equilibrium payoff vector $v = (v_1, v_2, v_3)$ must thus satisfy $v_1 + v_2 \geq 3$. In fact, it is possible to prove that the set of equilibrium payoffs of this repeated game is the set of feasible and individually rational payoffs that satisfy this constraint.

Approachability When the deviating player cannot be identified, it may be necessary to punish a group of suspects altogether. The notion of a payoff that is enforceable under group punishments is captured by the definition of approachable payoffs:

Definition 4 A payoff vector v is *approachable* if there exists a strategy profile σ such that, for every player i and unilateral deviation τ_i of player i , the average payoff of player i under (τ_i, σ_{-i}) is asymptotically less than or equal to v_i .

Blackwell's (1956) approachability theorem and its generalization by Kohlberg (1975) provide simple geometric characterizations of approachable payoffs. It is straightforward that approachability is a necessary condition on equilibrium payoffs:

Proposition 3 Every equilibrium payoff of the repeated game is approachable.

Renault and Tomala (2004) show that the conditions of Propositions 1 and 3 are tight for communication equilibria:

Theorem 10 For every game with imperfect monitoring, the set of communication equilibrium payoffs of the repeated game with no discounting is the set of approachable payoffs induced by distributions which are immune to undetectable deviations.

Tomala (1998) shows that pure strategy equilibrium payoffs of undiscounted repeated games with public signals are also characterized through identifiability and approachability conditions (the approachability definition then uses pure strategies). Tomala (1999) provides a similar characterization in mixed strategies for a restricted class of public signals.

Identification Through Endogenous Communication A deviation may be identified in the repeated game even though it cannot be identified in the stage game. In a *network game*, players are located at nodes of a graph, and each player monitors his neighbors' actions. Each player can use his actions as messages that are broadcasted to all the neighbors in the graph. The graph is called *2-connected* if no single node deletion disconnects the graph. Renault and Tomala (1998) show that when the graph is 2-connected, there exists a communication protocol among the players that ensures that the identity of any deviating player becomes common knowledge among all players in finite time. In this case, identification takes place through communication over the graph.

Public Equilibria

In a seminal paper, Green and Porter (1984) introduce a model in which firms are engaged in a production game and publicly observe market prices, which depend both on quantities produced and on non-observable exogenous market shocks. Can collusion be sustained at equilibrium even if prices convey imperfect information on quantities produced? This motivates the study of *public equilibria* for which sharp characterizations of equilibrium payoffs are obtained.

Signals are public when all sets of signals are identical, i.e., $Y_i = Y_{\text{pub}}$ for each i and $Q(\forall i, j, y_i = y_j | a) = 1$ for every a . A *public history* of length t is a record of t public signals, i.e., an element of $H_{\text{pub},t} = (Y_{\text{pub}})^t$. A strategy σ_i for player i is a *public strategy* if it depends on the public history only: if $h_i = (a_{i,1}, y_1, \dots, a_{i,t}, y_t)$ and $h'_i = (a'_{i,1}, y'_1, \dots, a'_{i,t}, y'_t)$ are two histories for player i such that $y_1 = y'_1, \dots, y_t = y'_t$, then $\sigma_i(h_i) = \sigma_i(h'_i)$.

Definition 5 A *perfect public equilibrium* is a profile of public strategies such that after every public history, each player's continuation strategy is a best reply to the opponents' continuation strategy profile.

The repetition of a Nash equilibrium of the stage game is a perfect public equilibrium, so that perfect public equilibria exist. Every perfect public equilibrium is a sequential equilibrium: any consistent belief assigns probability 1 to the realized public history and thus correctly forecasts future opponents' choices.

A Recursive Structure

A perfect public equilibrium (PPE henceforth) is a profile of public strategies that forms an equilibrium of the repeated game and such that, after every public history, the continuation strategy profile is also a PPE. The set of PPEs and of induced payoffs therefore possesses a recursive structure, as shown by Abreu et al. (1990). The argument is based on a dynamic programming principle. To state the main result, we first introduce some definitions.

Given a mapping $f : Y_{\text{pub}} \rightarrow \mathbb{R}^I$, $G(\delta, f)$ represents the one-shot game where each player i chooses actions in A_i and where payoffs are given by

$$(1 - \delta)g_i(a) + \delta \sum_{y \in Y_{\text{pub}}} Q(y|a)f_i(y).$$

In $G(\delta, f)$, the stage game is played, and players receive $f(y)$ as an additional payoff if y is the realized public signal. The weights $1 - \delta$ and δ

are the relative weights of present payoffs versus all future payoffs in the repeated game.

Definition 6 A payoff vector $u \in \mathbb{R}^I$ is decomposable with respect to the set $W \subset \mathbb{R}^I$ if there exists a mapping $f : Y_{\text{pub}} \rightarrow W$ such that v is a Nash equilibrium payoff of $G(\delta, f)$. $F_\delta(W)$ denotes the set of payoff vectors which are decomposable with respect to W .

Let $E(\delta)$ be the set of perfect public equilibrium payoffs of the repeated game discounted at the rate δ . The following result is due to Abreu et al. (1990):

Theorem 11 $E(\delta)$ is the largest bounded set W such that $W \subseteq F_\delta(W)$.

Fudenberg and Levine (1994) derive an asymptotic characterization of the set of PPE payoffs when the discount factor goes to 1 as follows. Given a vector $\lambda \in \mathbb{R}^I$, define the *score* in the direction λ as

$$k(\lambda) = \sup \langle \lambda, u \rangle$$

where the supremum is taken over the set of payoff vectors v that are Nash equilibrium payoffs of $G(\delta, f)$, where f is any mapping such that

$$\langle \lambda, u \rangle \geq \langle \lambda, f(y) \rangle, \quad \forall y \in Y_{\text{pub}}.$$

Scores are independent of the discount factor. The following theorem is due to Fudenberg and Levine (1994):

Theorem 12 Let C be the set of vectors v such that for every $\lambda \in \mathbb{R}^I$, $\langle \lambda, v \rangle \leq k(\lambda)$. If the interior of C is nonempty, $E(\delta)$ converges to C (for the Hausdorff topology) as δ goes to 1.

Fudenberg et al. (2007) relax the nonempty interior assumption. They provide an algorithm for computing the affine hull of $\lim_{\delta \rightarrow 1} E(\delta)$ and

provide a corresponding characterization of the set C with continuation payoffs belonging to this affine hull.

Folk Theorems for Public Equilibria

The recursive structure of Theorem 11 and the asymptotic characterization of PPE payoffs given by Theorem 12 are essential tools for finding sufficient conditions under which every feasible and individually rational payoff is an equilibrium payoff, i.e., conditions under which a Folk Theorem holds.

The two conditions under which a Folk Theorem in PPEs holds are (1) a condition of detectability of deviations and (2) a condition of identifiability of deviating players.

Definition 7 A profile of mixed actions $s = (s_i, s_{-i})$ has individual full rank if for each player i , the probability vectors (in the vector space $\mathbb{R}^{Y_{\text{pub}}}$)

$$\{Q(\cdot | a_i, s_{-i}) : a_i \in A_i\}$$

are linearly independent.

If s has individual full rank, no player can change the distribution of his actions without affecting the distribution of public signals. Individual full rank is thus a condition on detectability of deviations.

Definition 8 A profile of mixed actions s has pairwise full rank if for every pair of players $i \neq j$, the family of probability vectors

$$\{Q(\cdot | a_i, s_{-i}) : a_i \in A_i\} \cup \{Q(\cdot | a_j, s_{-j}) : a_j \in A_j\}$$

has rank $|A_i| + |A_j| - 1$.

Under the condition of pairwise full rank, deviations from two distinct players induce distinct distributions of public signals. Pairwise full rank is therefore a condition of identifiability of deviating players.

Fudenberg et al. (1994) prove the following theorem:

Theorem 13 Assume the set of feasible and individually rational payoff vectors F has nonempty interior. If every pure action profile has individual full rank and if there exists a mixed action profile

with pairwise full rank, then every convex and compact subset of the interior of F is a subset of $E(\delta)$ for δ large enough.

In particular, under the conditions of the theorem, every feasible and individually rational payoff vector is arbitrarily close to a PPE payoff for large enough discount factors. Variations of this result can be found in Fudenberg et al. (1994) and Fudenberg and Levine (1994).

Extensions

The Public Part of a Signal The definition of perfect public equilibria extends to the case in which each player's signal consists of two components: a public component and a private component. The public components of all players' signals are the same with probability 1. A public strategy is then a strategy that depends only on the public components of past signals, and all the analysis carries through.

Public Communication In the public communication extension of the repeated game, players make public announcements between any two stages of the repeated game. The profile of public announcements then forms a public signal, and recursive methods can be successfully applied. The fact that public communication is a powerful instrument to overcome the difficulties arising from private signals was first observed by Matsushima (1991a, b). Ben and Kahneman (1996), Kandori and Matsushima (1998), and Compte (1998) prove Folk Theorems in games with private signals and public communication. Kandori (2003) shows that in games with public monitoring, public communication allows to relax the conditions for the Folk Theorem of Fudenberg et al. (1994).

Private Strategies in Games with Public Monitoring PPE payoffs do not cover the full set of sequential equilibrium payoffs, even when signals are public, as some equilibria may rely on players using *private strategies*, i.e., strategies that depend on past chosen actions and past private signals. See Mailath et al. (2002) and Kandori

and Obara (2006) for examples. In a *minority game*, there are an odd number of players; each player chooses between actions A and B . Players choosing the least chosen (minority) action get a payoff of 1, and other players get 0. The public signal is the minority action. Renault et al. (2005, 2008) show that, for minority games, a Folk Theorem holds in private strategies but not in public strategies. Only few results are known concerning the set of sequential equilibrium payoffs in private strategies of games with public monitoring. A monotonicity property is obtained by Kandori (1992) who shows that the set of payoffs associated to sequential equilibria in pure strategies is increasing with respect to the quality of the public signal.

Almost-Public Monitoring Some PPEs are robust to small perturbations of public signals. Considering strategies with finite memory, Mailath and Morris (2002) identify a class of public strategies which are sequential equilibria of the repeated game with imperfect private monitoring, provided that the monitoring structure is close enough to a public one. They derive a Folk Theorem for games with almost-public and almost-perfect monitoring. Hörner and Olszewski (2007) strengthen this result and prove a Folk Theorem for games with almost-public monitoring. Under detectability and identifiability conditions, they prove that feasible and individually rational payoffs can be achieved by sequential equilibria with finite memory.

Almost-Perfect Monitoring

Monitoring is almost perfect when each player can identify the action profile of his opponents with near certainty. Almost-perfect monitoring is the natural framework to study the robustness of the Folk Theorem to small departures from the assumption that actions are perfectly observed.

The first results were obtained for the Prisoner's Dilemma. Sekiguchi (1997) shows that the cooperative outcome can be approximated at equilibrium when players are sufficiently patient and monitoring is almost perfect. Under the same assumptions, Bhaskar and Obara (2002), Piccione

(2002), and Ely and Valimaki (2002) show that a Folk Theorem obtains.

Piccione (2002) and Ely and Valimaki (2002) study a particular class of equilibria called *belief-free*. Strategies form a *belief-free* equilibrium if, whatever player i 's belief on the opponent's private history, the action prescribed by i 's strategy is a best response to the opponent's continuation strategy.

Ely et al. (2005) extend the belief-free approach to general games. However, they show that, in general, belief-free strategies are not enough to reconstruct a Folk Theorem, even when monitoring is almost perfect.

For general games and with any number of players, Hörner and Olszewski (2006) prove a Folk Theorem with almost-perfect monitoring. The strategies that implement the equilibrium payoffs are defined on successive blocks of a fixed length and are block belief-free in the sense that, at the beginning of every block, every player is indifferent between several continuation strategies, independently on his belief as to which continuation strategies are used by the opponents. This result closes the almost-perfect monitoring case by showing that equilibrium payoffs in the Folk Theorem are robust to a small amount of imperfect monitoring.

General Stochastic Signals

Besides the case of public (or almost-public) monitoring, little is known about equilibrium payoffs of repeated games with discounting and imperfect signals.

The Prisoner's Dilemma game is particularly important for economic applications. In particular, it captures the essential features of collusion with the possibility of secret price cutting, as in Stigler (1964).

When signals are imperfect, but independent conditionally on the pair of actions chosen (a condition called conditional independence), Matsushima (2004) shows that the efficient outcome of the repeated Prisoner's Dilemma game is an equilibrium outcome if players are sufficiently patient. In the equilibrium construction, every player's action is constant in every block. The conditional independence assumption is

crucial in that it implies that, during every block, a player has no feedback as to what signals the other player has received. The conditional independence assumption is nongeneric: it holds for a set of monitoring structures of empty interior.

Fong et al. (2007) prove that efficiency can be obtained at equilibrium without conditional independence. Their main assumption is that there exists a sufficiently informative signal, but this signal needs not be almost perfectly informative. Their result holds for a family of monitoring structures of nonempty interior. It is the first result that establishes cooperation in the Prisoner's Dilemma with impatient players for truly imperfect, private, and correlated signals.

Punishment Levels

Individual rationality is a key concept for Folk Theorems and equilibrium payoff characterizations. Given a repeated game, define the individually rational (IR) level of player i as the lowest payoff down to which this player may be punished in the repeated game.

Definition 9 The individual rational level of player i is

$$\lim_{\delta \rightarrow 1} \min_{\sigma_{-i}} \max_{\sigma_i} \mathbf{E}_{\sigma_i, \sigma_{-i}} \left[\sum_t (1 - \delta) \delta^{t-1} g_{i,t} \right]$$

where the min runs over profiles of behavior strategies for player $-i$ and the max over behavior strategies of player i .

That is, the individually rational level is the limit (as the discount factor goes to one) of the min max value of the discounted game (other approaches, through undiscounted games or limits of finitely repeated games, yield equivalent definitions; see Gossner and Tomala (2006)).

Comparison of the IR Level with the Min–Max
With perfect monitoring, the IR level of player i is player i 's min max in the one-shot game, as defined by Eq. 1. With imperfect monitoring, the IR level for player i is never larger than v_i since

player i 's opponents can force down player i to u_{-i} by repeatedly playing the min max strategy against player i .

With two players, it is a consequence of von Neumann's min max theorem (von Neumann 1928) that v_i is the IR level for player i .

For any number of players, Gossner and Hörner (2006) show that v_i is equal to the min max in the one-shot game whenever there exists a garbling of player i 's signal such that, conditionally on i 's garbled signal, the signals of i 's opponents are independent. Furthermore, the condition in Gossner and Hörner (2006) is also a necessary condition in normal form games extended by correlation devices (as in Aumann (1974)). A continuity result in the IR level also applies for monitoring structure close to those that satisfy the conditional independence condition.

The following example shows that, in general, the IR level can be lower than v_i :

Example 4 Consider the following three-player game. Player 1 chooses the row, player 2 the column, and player 3 the matrix. Players 1 and 2 perfectly observe the action profile, while player 3 observes player 2's action only. As we deal with the IR level of player 3, we specify the payoff for this player only.

		L	R		L	R
T	B	0	0		-1	0
B	W	0	-1		0	0
		W			E	

A simple computation shows that $u_3 = -\frac{1}{4}$ and that the min max strategies of players 1 and 2 are uniform. Consider the following strategies of players 1 and 2 in the repeated game: randomize uniformly at odd stages and play (T, L) or (B, R) depending on player 1's previous action at even stages. Against these strategies, player 3 cannot obtain better than $-\frac{1}{4}$ at odd stages and $-\frac{1}{2}$ at even stages, resulting in an average payoff of $-\frac{3}{8}$.

Entropy Characterizations

The exact computation of the IR level in games with imperfect monitoring requires to analyze the

optimal trade-off for punishing players between the production of correlated and private signals and the use of these signals for effective punishment. Gossner and Vieille (2002) and Gossner and Tomala (2006) develop tools based on information theory to analyze this trade-off. At any stage, the amount of correlation generated (or spent) by the punishing players is measured using the entropy function. Gossner and Tomala (2007) derive a characterization of the IR level for some classes of monitoring structures. Gossner et al. (2009) provide methods explicit computations of the IR level. In particular, for the above example, the IR level is computed and is about -0.401 . Explicit computations of IR levels for other games are derived by Goldberg (2007).

Acknowledgments The authors are grateful to Johannes Hörner for insightful comments.

Bibliography

Primary Literature

- Abreu D (1988) On the theory of infinitely repeated games with discounting. *Econometrica* 56:383–396
- Abreu D, Pearce D, Stacchetti E (1990) Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* 58:1041–1063
- Abreu D, Dutta P, Smith L (1994) The folk theorem for repeated games: a NEU condition. *Econometrica* 62:939–948
- Aumann RJ (1960) Acceptable points in games of perfect information. *Pac J Math* 10:381–417
- Aumann RJ (1964) Mixed and behavior strategies in infinite extensive games. In: Dresder M, Shapley LS, Tucker AW (eds) *Advances in game theory*. Princeton University Press, Princeton, pp 627–650
- Aumann RJ (1974) Subjectivity and correlation in randomized strategies. *J Math Econ* 1:67–95
- Aumann RJ (1981) Survey of repeated games. In: Aumann RJ (ed) *Essays in game theory and mathematical economics in honor of Oskar Morgenstern*. Wissenschaftsverlag, Bibliographisches Institut, Mannheim, pp 11–42
- Aumann RJ, Shapley LS (1976) Long-term competition – a game theoretic analysis. Re-edited in 1994. See Aumann and Shapley (1994)
- Aumann RJ, Shapley LS (1994) Long-term competition – a game theoretic analysis. In: Megiddo N (ed) *Essays on game theory*. Springer, New York, pp 1–15
- Ben EP, Kahneman M (1996) Communication in repeated games with private monitoring. *J Econ Theory* 70(2):281–297
- Benoît JP, Krishna V (1985) Finitely repeated games. *Econometrica* 53(4):905–922
- Bhaskar V, Obara I (2002) Belief-based equilibria in the repeated prisoners' dilemma with private monitoring. *J Econ Theory* 102:40–70
- Blackwell D (1951) Comparison of experiments. In: *Proceedings of the second Berkeley symposium on mathematical statistics and probability*. University of California Press, Berkeley, pp 93–102
- Blackwell D (1956) An analog of the minimax theorem for vector payoffs. *Pac J Math* 6:1–8
- Compte O (1998) Communication in repeated games with imperfect private monitoring. *Econometrica* 66:597–626
- Ely JC, Välimäki J (2002) A robust folk theorem for the prisoner's dilemma. *J Econ Theory* 102:84–106
- Ely JC, Hörner J, Olszewski W (2005) Belief-free equilibria in repeated games. *Econometrica* 73:377–415
- Fong K, Gossner O, Hörner J, Sannikov Y (2007) Efficiency in a repeated prisoner's dilemma with imperfect private monitoring. *Mimeo*
- Forges F (1986) An approach to communication equilibria. *Econometrica* 54:1375–1385
- Forges F, Mertens J-F, Neyman A (1986) A counterexample to the folk theorem with discounting. *Econ Lett* 20:7–7
- Friedman J (1971) A noncooperative equilibrium for supergames. *Rev Econ Stud* 38:1–12
- Fudenberg D, Levine DK (1994) Efficiency and observability with long-run and short-run players. *J Econ Theory* 62:103–135
- Fudenberg D, Maskin E (1986) The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* 54:533–554
- Fudenberg D, Levine DK, Maskin E (1994) The folk theorem with imperfect public information. *Econometrica* 62(5):997–1039
- Fudenberg D, Levine DK, Takahashi S (2007) Perfect public equilibrium when players are patient. *Game Econ Behav* 61:27–49
- Goldberg Y (2007) Secret correlation in repeated games with imperfect monitoring: the need for nonstationary strategies. *Math Oper Res* 32:425–435
- Gossner O (1995) The folk theorem for finitely repeated games with mixed strategies. *Int J Game Theory* 24:95–107
- Gossner O, Hörner J (2006) When is the individually rational payoff in a repeated game equal to the minmax payoff? DP 1440, CMS-EMS
- Gossner O, Tomala T (2006) Empirical distributions of beliefs under imperfect observation. *Math Oper Res* 31(1):13–30
- Gossner O, Tomala T (2007) Secret correlation in repeated games with signals. *Math Oper Res* 32:413–424
- Gossner O, Vieille N (2002) How to play with a biased coin? *Game Econ Behav* 41:206–226
- Gossner O, Laraki R, Tomala T (2009) Informationally optimal correlation. *Math Program B* 116:147–112
- Green EJ, Porter RH (1984) Noncooperative collusion under imperfect price information. *Econometrica* 52:87–100

- Hörner J, Olszewski W (2006) The folk theorem with private almost-perfect monitoring. *Econometrica* 74(6):1499–1544
- Hörner J, Olszewski W (2007) How robust is the folk theorem with imperfect public monitoring? Mimeo
- Kandori M (1992) The use of information in repeated games with imperfect monitoring. *Rev Econ Stud* 59:581–593
- Kandori M (2003) Randomization, communication, and efficiency in repeated games with imperfect public monitoring. *Econometrica* 71:345–353
- Kandori M, Matsushima H (1998) Private observation, communication and collusion. *Rev Econ Stud* 66:627–652
- Kandori M, Obara I (2006) Efficiency in repeated games revisited: the role of private strategies. *Econometrica* 74:499–519
- Kohlberg E (1975) Optimal strategies in repeated games with incomplete information. *Int J Game Theory* 4:7–24
- Kreps DM, Wilson RB (1982) Sequential equilibria. *Econometrica* 50:863–894
- Kuhn HW (1953) Extensive games and the problem of information. In: Kuhn HW, Tucker AW (eds) *Contributions to the theory of games*, vol II, vol 28, *Annals of mathematical studies*. Princeton University Press, Princeton, pp 193–216
- Lehrer E (1990) Nash equilibria of n -player repeated games with semi-standard information. *Int J Game Theory* 19:191–217
- Lehrer E (1991) Internal correlation in repeated games. *Int J Game Theory* 19:431–456
- Lehrer E (1992a) Correlated equilibria in two-player repeated games with nonobservable actions. *Math Oper Res* 17:175–199
- Lehrer E (1992b) Two players repeated games with non observable actions and observable payoffs. *Math Oper Res* 17:200–224
- Lehrer E, Pauzner A (1999) Repeated games with differential time preferences. *Econometrica* 67:393–412
- Luce RD, Raiffa H (1957) *Games and decisions: introduction and critical survey*. Wiley, New York
- Mailath G, Morris S (2002) Repeated games with almost-public monitoring. *J Econ Theory* 102:189–229
- Mailath GJ, Matthews SA, Sekiguchi T (2002) Private strategies in finitely repeated games with imperfect public monitoring. *Contrib Theor Econ* 2(1), Article 2
- Matsushima H (1991a) On the theory of repeated games with private information: part i: anti-folk theorem without communication. *Econ Lett* 35:253–256
- Matsushima H (1991b) On the theory of repeated games with private information: part ii: revelation through communication. *Econ Lett* 35:257–261
- Matsushima H (2004) Repeated games with private monitoring: two players. *Econometrica* 72:823–852
- Myerson RB (1982) Optimal coordination mechanisms in generalized principal-agent problems. *J Math Econ* 10:67–81
- Nash JF (1951) Noncooperative games. *Ann Math* 54:289–295
- Neyman A (1999) Cooperation in repeated games when the number of stages is not commonly known. *Econometrica* 67:45–64
- Piccione M (2002) The repeated prisoner's dilemma with imperfect private monitoring. *J Econ Theory* 102:70–84
- Renault J, Tomala T (1998) Repeated proximity games. *Int J Game Theory* 27:539–559
- Renault J, Tomala T (2004) Communication equilibrium payoffs of repeated games with imperfect monitoring. *Game Econ Behav* 49:313–344
- Renault J, Scarlatti S, Scarsini M (2005) A folk theorem for minority games. *Game Econ Behav* 53:208–230
- Renault J, Scarlatti S, Scarsini M (2008) Discounted and finitely repeated minority games with public signals. *Math Soc Sci* 56:44–74
- Rubinstein A (1977) *Equilibrium in supergames*. Center for Research in Mathematical Economics and Game Theory. Research memorandum, The Hebrew University, Jerusalem 25
- Rubinstein A (1979) Equilibrium in supergames with the overtaking criterion. *J Econ Theory* 21:1–9
- Rubinstein A (1994) Equilibrium in supergames. In: Megiddo N (ed) *Essays on game theory*. Springer, New York, pp 17–28
- Sekiguchi T (1997) Efficiency in repeated prisoner's dilemma with private monitoring. *J Econ Theory* 76:345–361
- Selten R (1965) Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit. *Z Gesamte Staatswiss* 12:201–324
- Smith L (1995) Necessary and sufficient conditions for the perfect finite horizon folk theorem. *Econometrica* 63:425–430
- Sorin S (1986) On repeated games with complete information. *Math Oper Res* 11:147–160
- Stigler G (1964) A theory of oligopoly. *J Polit Econ* 72:44–61
- Tomala T (1998) Pure equilibria of repeated games with public observation. *Int J Game Theory* 27:93–109
- Tomala T (1999) Nash equilibria of repeated games with observable payoff vector. *Game Econ Behav* 28:310–324
- von Neumann J (1928) Zur Theorie der Gesellschaftsspiele. *Math Ann* 100:295–320
- Wen Q (1994) The “folk theorem” for repeated games with complete information. *Econometrica* 62:949–954

Books and Reviews

- Mailath GJ, Samuelson L (2006) *Repeated games and reputations: long-run relationships*. Oxford University Press, Oxford
- Mertens J-F (1986) Repeated games. In: *Proceedings of the international congress of mathematicians*, Berkeley, pp 1528–1577
- Mertens J-F, Sorin S, Zamir S (1994) *Repeated games*. CORE discussion paper. Université Catholique de Louvain, Louvain-la-neuve, pp 9420–9422

Repeated Games with Incomplete Information

Jérôme Renault
Toulouse School of Economics, Université
Toulouse 1 Capitole, Toulouse, France

Article Outline

Glossary and Notation
Definition of the Subject and Its Importance
Strategies, Payoffs, Value, and Equilibria
The Standard Model of Aumann and Maschler
Vector Payoffs and Approachability
Zero-Sum Games with Lack of Information on
Both Sides
Nonzero-sum Games with Lack of Information on
One Side
Nonobservable Actions
Advances
Future Directions
Bibliography

Glossary and Notation

Repeated game with incomplete information A situation where several players repeat the same stage game, the players having different knowledge of the stage game which is repeated.

Strategy of a player A rule, or program, describing the action taken by the player in any possible case which may happen.

Strategy profile A vector containing a strategy for each player.

Lack of information on one side Particular case where all the players but one perfectly know the stage game which is repeated.

Zero-sum games 2-player games where the players have opposite payoffs.

Value Solution (or price) of a zero-sum game, in the sense of the fair amount that player 1 should give to player 2 to be entitled to play the game.

Equilibrium Strategy profile where each player's strategy is in best reply against the strategy of the other players.

Completely revealing strategy Strategy of a player which eventually reveals to the other players everything known by this player on the selected state.

Non revealing strategy Strategy of a player which reveals nothing on the selected state.

The simplex of probabilities over a finite set For a finite set S , we denote by $\Delta(S)$ the set of probabilities over S , and we identify $\Delta(S)$ to $\{p = (p_s)_{s \in S} \in \mathbb{R}^S, \forall s \in S p_s \geq 0 \text{ and } \sum_{s \in S} p_s = 1\}$. Given s in S , the Dirac measure on s will be denoted by δ_s . For $p = (p_s)_{s \in S}$ and $q = (q_s)_{s \in S}$ in $\mathbb{R}^S$, we will use, unless otherwise specified, $\|p - q\| = \sum_{s \in S} |p_s - q_s|$.

Definition of the Subject and Its Importance

Introduction

In a repeated game with incomplete information, there is a basic interaction called stage game which is repeated over and over by several participants called players. The point is that the players do not perfectly know the stage game which is repeated, but rather have different knowledge about it. As illustrative examples, one may think of the following situations: an oligopolistic competition where firms do not know the production costs of their opponents, a financial market where traders bargain over units of an asset which terminal value is imperfectly known, a cryptographic model where some participants want to transmit some information (e.g., a credit card number) without being understood by other participants, a conflict when a particular side may be able to understand the communications inside the opponent side (or might have a particular type of weapons),...

Natural questions arising in this context are as follows. What is the optimal behavior of a player with a perfect knowledge of the stage game? Can we

determine which part of the information such a player should use? Can we price the value of possessing a particular information? How should one player behave while having only a partial information?

Foundations of games with incomplete information have been studied in (Harsanyi 1967; Mertens and Zamir 1985). Repeated games with incomplete information have been introduced in the sixties by Aumann and Maschler (1995), and we present here the basic and fundamental results of the domain. Let us start with a few well-known elementary examples (Aumann and Maschler 1995; Zamir 1992).

Basic Examples In each example, there are two players, and the game is zero-sum, i.e., player 2's payoff always is the opposite of player 1's payoff. There are two states a and b , and the possible stage games are given by two real matrices G^a and G^b with identical size. Initially a true state of nature $k \in \{a, b\}$ is selected with even probability between a and b , and k is announced to player 1 only. Then the matrix game G^k is repeated over and over: at every stage, simultaneously player 1 chooses a row i , whereas player 2 chooses a column j , the stage payoff for player 1 is then $G^k(i, j)$, but only i and j are publicly announced before proceeding to the next stage. Players are patient and want to maximize their long-run average expected payoffs.

Example 1 $G^a = \begin{pmatrix} 0 & 0 \\ 0 & -1 \end{pmatrix}$ and $G^b = \begin{pmatrix} -1 & 0 \\ 0 & 0 \end{pmatrix}$.

This example is trivial. In order to maximize his payoff, player 1 just has to play, at any stage, the *Top* row if the state is a and the *Bottom* row if the state is b .

Example 2 $G^a = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$ and $G^b = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$.

A naive strategy for player 1 would be to play at stage 1: *Top* if the state is a , and *Bottom* if the state is b . Such a strategy is called completely revealing, or CR, because it allows player 2 to deduce the selected state from the observation of the actions played by player 1. This strategy of player 1 would be optimal here if a single stage was to be played, but it is a very weak strategy on the long run and does not guarantee more than

zero at each stage $t \geq 2$ (because player 2 can play *Left* or *Right* depending on player 1's first action).

On the opposite, player 1 may not use his information and play a nonrevealing, or NR, strategy, i.e., a strategy which is independent of the selected state.

He can consider the average matrix $\frac{1}{2}G^a + \frac{1}{2}G^b = \begin{pmatrix} 1/2 & 0 \\ 0 & 1/2 \end{pmatrix}$ and play independently at each stage an optimal mixed action in this matrix, i.e., here the unique mixed action $\frac{1}{2}\text{Top} + \frac{1}{2}\text{Bottom}$. It will turn out that this is here the optimal behavior for player 1, and the value of the repeated game is the value of the average matrix, i.e., $1/4$.

Example 3 $G^a = \begin{pmatrix} 4 & 0 & 2 \\ 4 & 0 & -2 \end{pmatrix}$ and $G^b = \begin{pmatrix} 0 & 4 & -2 \\ 0 & 4 & 2 \end{pmatrix}$.

Playing a CR strategy for player 1 does not guarantee more than zero in the long-run, because player 2 will eventually be able to play *Middle* if the state is a , and *Left* if the state is b . But a NR strategy will not do better, because the average matrix $\frac{1}{2}G^a + \frac{1}{2}G^b$ is $\begin{pmatrix} 2 & 2 & 0 \\ 2 & 2 & 0 \end{pmatrix}$, hence has value 0.

We will see later that an optimal strategy for player 1 in this game is to play as follows. Initially, player 1 chooses an element s in $\{T, B\}$ as follows: if $k = a$, then $s = T$ with probability $3/4$, and thus $s = B$ with probability $1/4$; and if $k = b$, then $s = T$ with probability $1/4$, and $s = B$ with probability $3/4$. Then at each stage player 1 plays row s , independently of the actions taken by player 2. The conditional probabilities satisfy: $P(k = a | s = T) = 3/4$, and $P(k = a | s = B) = 1/4$. At the end of stage 1, player 2 will have learnt, from the action played by his opponent, something about the selected state: his belief on the state will move from $\frac{1}{2}a + \frac{1}{2}b$ to $\frac{3}{4}a + \frac{1}{4}b$ or to $\frac{1}{4}a + \frac{3}{4}b$. But player 2 still does not know perfectly the selected state. Such a strategy of player 1 is called *partially revealing*.

General Definition

Formally, a repeated game with incomplete information is given by the following data. There is a set of players N and a set of states K . Each player i in N has a set of actions A^i and a set of signals U^i ,

and we denote by $A = \prod_{i \in N} A^i$ the set of action profiles and by $U = \prod_{i \in N} U^i$ the set of signal profiles. Every player i has a payoff function $g^i: K \times A \rightarrow \mathbb{R}$. There is a signaling function $q: K \times A \rightarrow \Delta(U)$, and an initial probability $\pi \in \Delta(K \times U)$. In what follows, we will always assume the sets of players, states, actions, and signals to be nonempty and finite.

A repeated game with incomplete information can thus be denoted by $\Gamma = (N, K, (A^i)_{i \in N}, (U^i)_{i \in N}, (g^i)_{i \in N}, q, \pi)$. The progress of the game is the following.

- Initially, an element $(k, (u_0^i)_i)$ is selected according to π : k is the realized state of nature and will remain fixed, and each player i learns u_0^i (and nothing more than u_0^i).
- At each integer stage $t \geq 1$, simultaneously every player i chooses an action a_t^i in A^i , and we denote by $a_t = (a_t^i)_i$ the action profile played at stage t . The stage payoff of a player i is then given by $g^i(k, a_t)$. A signal profile $(u_t^i)_i$ is selected according to $q(k, a_t)$, and each player i learns u_t^i (and nothing more than u_t^i) before proceeding to the next stage.

Remarks

- The players do not necessarily know their stage payoff after each stage (as an illustration, imagine the players bargaining over units of an asset which terminal value will only be known “at the end” of the game). This is without loss of generality, because it is possible to add hypotheses on q so that each player will be able to deduce his stage payoff from his realized stage signal.
- Repeated games with complete information are a particular case, corresponding to the situation where each initial signal u_0^i reveals the selected state. Such games are studied in the chapter ▶ “Repeated Games with Complete Information”
- Games where the state variable k evolve from stage to stage, according to the actions played, are called stochastic games. These games are not covered here, but in a specific chapter entitled ▶ “Stochastic Games”.

- The most standard case of signaling function is when each player exactly learns, at the end of each stage t , the whole action profile a_t . Such games are usually called games with “perfect monitoring,” “full monitoring,” “perfect observation” or with “observable actions.”

Strategies, Payoffs, Value, and Equilibria

Strategies

A (behavior) strategy for player i is a rule, or program, describing the action taken by this player in any possible case which may happen. These actions may be chosen at random, so a strategy for player i is an element $\sigma^i = (\sigma_t^i)_{t \geq 1}$, where for each t , σ_t^i is a mapping from $U^i \times (U^i \times A^i)^{t-1}$ to $\Delta(A^i)$ giving the lottery played by player i at stage t as a function of the past signals and actions of player i . The set of strategies for player i is denoted by Σ^i .

A history of length t in Γ is a sequence $(k, u_0, a_1, u_1, \dots, a_t, u_t)$, and the set of such histories is the finite set $K \times U \times (A \times U)^t$. An infinite history is called a play, and the set of plays is denoted by $\Omega = K \times U \times (A \times U)^\infty$ and is endowed with the product σ -algebra. A strategy profile $\sigma = (\sigma^i)_i$ naturally induces, together with the initial probability π , a probability distribution over the set of histories of length t . This probability uniquely extends to a probability over plays and is denoted by $\mathbb{P}_{\pi, \sigma}$.

Payoffs

Given a time horizon T , the average expected payoff of player i , up to stage T , if the strategy profile σ is played, is denoted by:

$$\gamma_T^i(\sigma) = \mathbb{E}_{\mathbb{P}_{\pi, \sigma}} \left(\frac{1}{T} \sum_{t=1}^T g^i(k, a_t) \right).$$

The T -stage game is the game Γ_T where simultaneously each player i chooses a strategy σ^i in Σ^i , then receives the payoff $\gamma_T^i((\sigma^j)_{j \in N})$.

Given a discount factor λ in $(0, 1]$, the λ -discounted payoff of player i is denoted by:

$$\gamma_\lambda^i(\sigma) = \mathbb{E}_{\mathbb{P}_{\pi, \sigma}} \left(\lambda \sum_{t=1}^{\infty} (1 - \lambda)^{t-1} g^i(k, a_t) \right).$$

The λ -discounted game is the game Γ_λ where simultaneously, each player i chooses a strategy σ^i in Σ^i , then receives the payoff $\gamma_\lambda^i((\sigma^j)_{j \in N})$.

Remark A strategy for player i is called *pure* if it always plays in a deterministic way. A *mixed strategy* for player i is defined as a probability distribution over the set of pure strategies (endowed with the product σ -algebra). Kuhn's theorem (see Aumann (1964), Kuhn (1953) or Sorin (2002) for a modern presentation) states that mixed strategies or behavior strategies are equivalent, in the following sense: for each behavior strategy σ^i , there exists a mixed strategy τ^i of the same player such that $\mathbb{P}_{\pi, \sigma^i, \sigma^{-i}} = \mathbb{P}_{\pi, \tau^i, \sigma^{-i}}$ for any strategy profile σ^{-i} of the other players, and vice versa if we exchange the words “behavior” and “mixed.” Unless otherwise specified, the word strategy will refer here to a behavior strategy, but we will also sometimes equivalently use mixed strategies, or even mixtures of behavior strategies.

Value of Zero-Sum Games

By definition the game is zero-sum if there are two players, say player 1 and player 2, with opposite payoffs. The T -stage game Γ_T can then be seen as a matrix game; hence, by the minmax theorem it has a value $v_T = \sup_{\sigma^1} \inf_{\sigma^2} \gamma_T^1(\sigma^1, \sigma^2) = \inf_{\sigma^2} \sup_{\sigma^1} \gamma_T^1(\sigma^1, \sigma^2)$. Similarly, one can use Sion's theorem (1958) to show that the λ -discounted game has a value $v_\lambda = \sup_{\sigma^1} \inf_{\sigma^2} \gamma_\lambda^1(\sigma^1, \sigma^2) = \inf_{\sigma^2} \sup_{\sigma^1} \gamma_\lambda^1(\sigma^1, \sigma^2)$.

To study long term strategic aspects, it is also important to consider the following notion of uniform value. Players are asked to play well uniformly in the time horizon, i.e., simultaneously in all game Γ_T with T sufficiently large (or similarly uniformly in the discount factor, i.e., simultaneously in all game Γ_λ with λ sufficiently low).

Definitions 1 Player 1 can guarantee the real number v in the repeated game Γ if: $\forall \varepsilon > 0$, $\exists \sigma^1 \in \Sigma^1$, $\exists T_0$, $\forall T \geq T_0$, $\forall \sigma^2 \in \Sigma^2$, $\gamma_T^1(\sigma^1, \sigma^2) \geq v - \varepsilon$. Similarly, Player 2 can guarantee v in Γ if $\forall \varepsilon > 0$, $\exists \sigma^2 \in \Sigma^2$, $\exists T_0$, $\forall T \geq T_0$, $\forall \sigma^1 \in \Sigma^1$, $\gamma_T^1(\sigma^1, \sigma^2) \leq v + \varepsilon$. If both player 1 and player 2 can guarantee v , then v is called the uniform value of the repeated game. A strategy σ^1 of player 1 satisfying $\exists T_0$, $\forall T \geq T_0$, $\forall \sigma^2 \in \Sigma^2$, $\gamma_T^1(\sigma^1, \sigma^2) \geq v$ is then called an optimal strategy of player 1 (optimal strategies of player 2 are defined similarly).

The uniform value, whenever it exists, is necessarily unique. Its existence is a strong property, which implies that both v_T as T goes to infinity, and v_λ as λ goes to zero, converge to the uniform value.

Equilibria of General-Sum Games

In the general case, the T -stage game Γ_T can be seen as the mixed extension of a finite game and consequently possesses a Nash equilibrium. Similarly, the discounted game Γ_λ always has, by the Nash Glicksberg theorem, a Nash equilibrium. Concerning uniform notions, couples of optimal strategies are generalized as follows.


Definitions 2 A strategy profile $\sigma = (\sigma^i)_{i \in N}$ is a uniform Nash equilibrium of Γ if: (1) $\forall \varepsilon > 0$, σ is an ε -Nash equilibrium in every finitely repeated game sufficiently long, that is, $\exists T_0$, $\forall T \geq T_0$, $\forall i \in N$, $\forall \tau^i \in \Sigma^i$, $\gamma_T^i(\tau^i, \sigma^{-i}) \leq \gamma_T^i(\sigma) + \varepsilon$, and (2) the sequence of payoffs $\left((\gamma_T^i(\sigma))_{i \in N} \right)_T$ converges to a limit payoff $(\gamma^i(\sigma))_{i \in N}$ in $\mathbb{R}^N$.

Remark The initial probability π will play a great role in the following analyses, so we will often write $\gamma_T^{i, \pi}(\sigma)$ for $\gamma_T^i(\sigma)$, $v_T(\pi)$ for the value v_T , etc. . .

The Standard Model of Aumann and Maschler

This famous model has been introduced in the sixties by Aumann and Maschler (see the reedition (Aumann and Maschler 1995)). It deals with zero-sum games with lack of information on one side and observable actions, as in the basic examples previously presented. There is a finite set of states K , an initial probability $p = (p^k)_{k \in K}$ on K , and a family of matrix games G^k with identical

size $I \times J$. Initially, a state k in K is selected according to p and announced to player 1 (called the informed player) only. Then the matrix game G^k is repeated over and over: at every stage, simultaneously player 1 chooses a row i in I , whereas player 2 chooses a column j in J , the stage payoff for player 1 is then $G^k(i, j)$, but only i and j are publicly announced before proceeding to the next stage. Denote by M the constant $\max_{k,i,j} |G^k(i, j)|$.



Repeated Games with Incomplete Information, Fig. 1 Splitting

Basic Tools: Splitting, Martingale, Concavification, and the Recursive Formula

The following aspects are simple but fundamental. The initial probability $p = (p^k)_{k \in K}$ represents the initial belief, or *a priori*, of player 2 on the selected state of nature. Assume that player 1 chooses his first action (or more generally a message or signal s from a finite set S) according to a probability distribution depending on the state, i.e., according to a transition probability $x = (x^k)_{k \in K} \in \Delta(S)^K$. For each signal s , the probability that s is chosen is denoted $\lambda(x, s) = \sum_k p^k x^k(s)$, and given s such that $\lambda(x, s) > 0$ the conditional probability on K , or *a posteriori* of player 2, is $\hat{p}(x, s) = \left(\frac{p^k x^k(s)}{\lambda(x, s)} \right)_{k \in K}$.

We clearly have:

$$p = \sum_{s \in S} \lambda(x, s) \hat{p}(x, s). \quad (1)$$

So the *a priori* p lies in the convex hull of the *a posteriori*. The following lemma expresses a reciprocal: player 1 is able to induce any family of *a posteriori* containing p in its convex hull.

Splitting Lemma 1 Assume that p is written as a convex combination $p = \sum_{s \in S} \lambda_s p_s$ with positive coefficients. There exists a transition probability $x \in \Delta(S)^K$ such that $\forall s \in S$, $\lambda_s = \lambda(x, s)$ and $p_s = \hat{p}(x, s)$.

Proof Just put $x^k(s) = \frac{\lambda_s p_s^k}{p^k}$ if $p^k > 0$. (Fig. 1)

Equation 1 not only tells that the *a posteriori* contains p in their convex hull, but also that the expectation of the *a posteriori* is the *a priori*. We

are here in a repeated context, and for every strategy profile σ one can define the process $(p_t(\sigma))_{t \geq 0}$ of the *a posteriori* of player 2. We have $p_0 = p$, and $p_t(\sigma)$ is the random variable of player 2's belief on the state after the first t stages. More precisely, we define for any $t \geq 0$, $h_t = (i_1, j_1, \dots, i_t, j_t) \in (I \times J)^t$ and $k \in K$:

$$p_t^k(\sigma, h_t) = \mathbb{P}_{p, \sigma}(k | h_t) = \frac{p^k \mathbb{P}_{\delta^k, \sigma}(h_t)}{\mathbb{P}_{p, \sigma}(h_t)}.$$

$p_t(\sigma, h_t) = (p_t^k(\sigma, h_t))_{k \in K} \in \Delta(K)$ (arbitrarily defined if $\mathbb{P}_{p, \sigma}(h_t) = 0$) is the conditional probability on the state of nature given that σ is played and h_t has occurred in the first t stages. It is easy to see that as soon as $\mathbb{P}_{p, \sigma}(h_t) > 0$, $p_t(\sigma, h_t)$ does not depend on player 2's strategy σ^2 , nor on player 2's last action j_t . It is fundamental to notice that:

Martingale of a Posteriori Lemma 2 $(p_t(\sigma))_{t \geq 0}$ is a $\mathbb{P}_{p, \sigma}$ -martingale with values in $\Delta(K)$.

This is indeed a general property of Bayesian learning of a fixed unknown parameter: the expectation of what I will know tomorrow is what I know today. This martingale is controlled by the informed player, and the splitting lemma shows that this player can essentially induce any martingale issued from the *a priori* p . Notice that to be able to compute the realizations of the martingale, player 2 needs to know the strategy σ^1 used by player 1.

The splitting lemma also easily gives the following concavification result. Let f be a continuous mapping from $\Delta(K)$ to $\mathbb{R}$. The smallest concave function above f is denoted by $\text{cav } f$, and we have:

$$\begin{aligned} \text{cav} f(p) &= \max \left\{ \sum_{s \in S} \lambda_s f(p_s), S \text{ finite}, \forall s \right. \\ &\quad \left. \lambda_s \geq 0, p_s \in \Delta(K), \sum_{s \in S} \lambda_s \right. \\ &\quad \left. = 1, \sum_{s \in S} \lambda_s p_s = p \right\} \end{aligned}$$

Concavification Lemma 3 If for any initial probability p , the informed player can guarantee $f(p)$ in the game $\Gamma(p)$, then for any p this player can also guarantee $\text{cav} f(p)$ in $\Gamma(p)$.

Nonrevealing Games


As soon as player 1 uses a strategy which depends on the selected state, the martingale of a posteriori will move and player 2 will have learnt something on the state. This is the dilemma of the informed player: he cannot use the information on the state without revealing information. Imagine now that player 1 decides to reveal no information on the selected state and plays independently of it. Since payoffs are defined via expectations, it is as if the players were repeating the average matrix game $G(p) = \sum_{k \in K} p^k G^k$. Its value is:

$$\begin{aligned} u(p) &= \max_{x \in \Delta(I)} \min_{y \in \Delta(J)} \sum_{i,j} x(i)y(j)G(p)(i,j) \\ &= \min_{y \in \Delta(J)} \max_{x \in \Delta(I)} \sum_{i,j} x(i)y(j)G(p)(i,j). \end{aligned}$$

u is a Lipschitz function, with constant M , from $\Delta(K)$ to $\mathbb{R}$. Clearly, player 1 can guarantee $u(p)$ in the game $\Gamma(p)$ by playing i.i.d. at each stage an optimal strategy in $G(p)$. By the concavification lemma, we obtain:

Proposition 1 Player 1 can guarantee $\text{cav} u(p)$ in the game $\Gamma(p)$.

Let us come back to the examples. In Example 1, we have $u(p) = \text{Val} \begin{pmatrix} -(1-p) & 0 \\ 0 & -p \end{pmatrix} = -p(1-p)$, where $p \in [0, 1]$ stands here for the probability of state a . This is a convex function of p , and $\text{cav} u(p) = 0$ for all p . In Example 2, $u(p) = p(1-p)$ for all p ; hence, u is already concave and $\text{cav} u = u$. Regarding Example 3, the following picture shows the functions u (regular line) and $\text{cav} u$ (dashed line) (Fig. 2).



Repeated Games with Incomplete Information, Fig. 2 u and $\text{cav} u$

Let us consider again the partially revealing strategy previously described. With probability $1/2$, the a posteriori will be $\frac{3}{4}a + \frac{1}{4}b$, and player 1 will play *Top* which is optimal in $\frac{3}{4}G^a + \frac{1}{4}G^b = \begin{pmatrix} 3 & 1 & 1 \\ 3 & 1 & -1 \end{pmatrix}$. Similarly with probability $1/2$, the a posteriori will be $\frac{1}{4}a + \frac{3}{4}b$ and player 1 will play an optimal strategy in $\frac{1}{4}G^a + \frac{3}{4}G^b$. Consequently, this strategy guarantees $1/2 u(3/4) + 1/2 u(1/4) = \text{cav} u(1/2) = 1$ to player 1.

Player 2 Can Guarantee the Limit Value

In the infinitely repeated game with initial probability p , player 2 can play as follows: T being fixed, he can play an optimal strategy in the T -stage game $\Gamma_T(p)$, then forget everything and play again an optimal strategy in the T -stage game $\Gamma_T(p)$, etc. By doing so, he guarantees $v_T(p)$ in the game $\Gamma(p)$. So he can guarantee $\inf_T v_T(p)$ in this game, and this implies that $\limsup_T v_T(p) \leq \inf_T v_T(p)$. As a consequence, we obtain:

Proposition 2 The sequence $(v_T(p))_T$ converges to $\inf_T v_T(p)$, and this limit can be guaranteed by player 2 in the game $\Gamma(p)$.

Uniform Value: cavu Theorem

We will see here that $\lim_T v_T(p)$ is nothing but $\text{cav} u(p)$, and since this quantity can be guaranteed by both players, this is the uniform value of the game $\Gamma(p)$. The idea of the proof is the following. The martingale $(p_i(\sigma))_i \geq 0$ is bounded, hence will converge almost surely, and we have a bound on its L^1 variation (see Lemma 4 below). This means that after a certain stage the martingale will essentially remain constant, so approximately player 1 will play in a nonrevealing way, so will not be able to have a stage payoff greater than $u(q)$, where q is a “limit a posteriori.” Since the expectation of the a posteriori is the

a priori p , player 1 cannot guarantee more than $\max \{\sum_{s \in S} \lambda_s u(p_s), S \text{ finite}, \forall s \in S \lambda_s \geq 0, p_s \in \Delta(K), \sum_{s \in S} \lambda_s = 1, \sum_{s \in S} \lambda_s p_s = p\}$, that is, more than $\text{cav}u(p)$. Let us now proceed to the formal proof.

Fix a strategy σ^1 of player 1, and define the strategy σ^2 of player 2 as follows: play at each stage an optimal strategy in the matrix game $G(p_t)$, where p_t is the current *a posteriori* in $\Delta(K)$. Assume that $\sigma = (\sigma^1, \sigma^2)$ is played in the repeated game $\Gamma(p)$. To simplify notations, we write $\mathbb{P}$ for $\mathbb{P}_{p, \sigma, p_t}(h_t)$ for $p_t(\sigma, h_t)$, etc. We use everywhere norms $\|\cdot\|$. To avoid confusion between variables and random variables in the following computations, we will use tildes to denote random variables, e.g., $\tilde{k}$ will denote the random variable of the selected state.

Lemma 4

$$\forall T \geq 1, \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}(\|p_{t+1} - p_t\|) \leq \frac{\sum_{k \in K} \sqrt{p^k(1-p^k)}}{\sqrt{T}}.$$

Proof This is a property of martingales with values in $\Delta(K)$ and expectation p . We have for each state k and $t \geq 0$: $\mathbb{E}((p_{t+1}^k - p_t^k)^2) = \mathbb{E}(\mathbb{E}((p_{t+1}^k - p_t^k)^2 | \mathcal{H}_t))$, where $\mathcal{H}_t$ is the σ -algebra on plays generated by the first t action profiles. So $\mathbb{E}((p_{t+1}^k - p_t^k)^2) = \mathbb{E}(\mathbb{E}((p_{t+1}^k)^2 + (p_t^k)^2 - 2p_{t+1}^k p_t^k | \mathcal{H}_t)) = \mathbb{E}((p_{t+1}^k)^2) - \mathbb{E}((p_t^k)^2)$. So $\mathbb{E}(\sum_{t=0}^{T-1} (p_{t+1}^k - p_t^k)^2) = \mathbb{E}((p_T^k)^2) - (p^k)^2 \leq p^k(1-p^k)$. By Cauchy-Schwartz inequality, we also have for each k ,

$$\mathbb{E}(\frac{1}{T} \sum_{t=0}^{T-1} |p_{t+1}^k - p_t^k|) \leq \sqrt{\frac{1}{T} \mathbb{E}(\sum_{t=0}^{T-1} (p_{t+1}^k - p_t^k)^2)}$$

and the result follows. $\square$

For h_t in $(I \times J)^t$, $\sigma_{t+1}^1(k, h_t)$ is the mixed action in $\Delta(I)$ played by player 1 at stage $t+1$ if the state is k and h_t has previously occurred, and we write $\bar{\sigma}_{t+1}^1(h_t)$ for the law of the action of player 1 of stage $t+1$ after h_t : $\bar{\sigma}_{t+1}^1(h_t) = \sum_{k \in K} p_t^k(h_t) \sigma_{t+1}^1(k, h_t) \in \Delta(I)$. $\bar{\sigma}_{t+1}^1(h_t)$ can be seen as the average action played by player 1 after h_t and will be used as a nonrevealing approximation for $(\sigma_{t+1}^1(k, h_t))_k$. The next lemma precisely links the variation of the martingale $(p_t(\sigma))_{t \geq 0}$, i.e., the information revealed by player 1, and the dependence of player 1's action on the selected state, i.e., the information used by player 1.

Lemma 5

$$\begin{aligned} \forall t \geq 0, \forall h_t \in (I \times J)^t, \mathbb{E}(\|p_{t+1} - p_t\| | h_t) \\ = \mathbb{E}(\|\sigma_{t+1}^{\tilde{k}}(h_t) - \bar{\sigma}_{t+1}^1(h_t)\| | h_t). \end{aligned}$$

Proof Fix $t \geq 0$ and h_t in $(I \times J)^t$ s.t. $\mathbb{P}_{p, \sigma}(h_t) > 0$. For (i_{t+1}, j_{t+1}) in $I \times J$, one has:

$$\begin{aligned} p_{t+1}^k(h_t, i_{t+1}, j_{t+1}) &= \mathbb{P}(\tilde{k} = k | h_t, i_{t+1}) \\ &= \frac{\mathbb{P}(\tilde{k} = k | h_t) \mathbb{P}(i_{t+1} | k, h_t)}{\mathbb{P}(i_{t+1} | h_t)} \\ &= \frac{p_t^k(h_t) \sigma_{t+1}^1(k, h_t)(i_{t+1})}{\bar{\sigma}_{t+1}^1(h_t)(i_{t+1})}. \end{aligned}$$

Consequently,

$$\begin{aligned} \mathbb{E}(\|p_{t+1} - p_t\| | h_t) &= \sum_{i_{t+1} \in I} \bar{\sigma}_{t+1}^1(h_t)(i_{t+1}) \\ &\quad \sum_{k \in K} |p_{t+1}^k(h_t, i_{t+1}) - p_t^k(h_t)| \\ &= \sum_{i_{t+1} \in I} \sum_{k \in K} |p_t^k(h_t) \sigma_{t+1}^1(k, h_t)(i_{t+1}) \\ &\quad - \bar{\sigma}_{t+1}^1(h_t)(i_{t+1}) p_t^k(h_t)| \\ &= \sum_{k \in K} p_t^k(h_t) \|\sigma_{t+1}^1(k, h_t) - \bar{\sigma}_{t+1}^1(h_t)\| \\ &= \mathbb{E}(\|\sigma_{t+1}^{\tilde{k}}(\tilde{k}, h_t) - \bar{\sigma}_{t+1}^1(h_t)\| | h_t). \end{aligned}$$

We can now control payoffs. For $t \geq 0$ and h_t in $(I \times J)^t$:

$$\begin{aligned} \mathbb{E}(G^{\tilde{k}}(\tilde{i}_{t+1}, \tilde{j}_{t+1}) | h_t) \\ &= \sum_{k \in K} p_t^k(h_t) G^k(\sigma_{t+1}^1(k, h_t), \sigma_{t+1}^2(h_t)) \\ &\leq \sum_{k \in K} p_t^k(h_t) G^k(\bar{\sigma}_{t+1}^1(h_t), \sigma_{t+1}^2(h_t)) \\ &\quad + M \sum_{k \in K} p_t^k(h_t) \|\sigma_{t+1}^1(k, h_t) \\ &\quad - \bar{\sigma}_{t+1}^1(h_t)\| \leq u(p_t(h_t)) \\ &\quad + M \sum_{k \in K} p_t^k(h_t) \|\sigma_{t+1}^1(k, h_t) - \bar{\sigma}_{t+1}^1(h_t)\|, \end{aligned}$$

where $u(p_t(h_t))$ comes from the definition of σ^2 . By Lemma 5, we get:

$$\begin{aligned} \mathbb{E}(G^{\tilde{k}}(\tilde{i}_{t+1}, \tilde{j}_{t+1}) | h_t) &\leq u(p_t(h_t)) \\ &\quad + M \mathbb{E}(\|p_{t+1} - p_t\| | h_t). \end{aligned}$$

Applying Jensen's inequality yields:

$$\mathbb{E}\left(G^{\tilde{k}}(\tilde{i}_{t+1}, \tilde{j}_{t+1})\right) \leq \text{cavu}(p) + M \mathbb{E}(\|p_{t+1} - p_t\|).$$

We now apply Lemma 4 and obtain:

$$\begin{aligned} \gamma_T^{1,p}(\sigma^1, \sigma^2) &= \mathbb{E}\left(\frac{1}{T} \sum_{t=0}^{T-1} G^{\tilde{k}}(\tilde{i}_{t+1}, \tilde{j}_{t+1})\right) \\ &\leq \text{cavu}(p) + \frac{M}{\sqrt{T}} \sum_{k \in K} \sqrt{p^k(1-p^k)}. \end{aligned}$$

This is true for any strategy σ^1 of player 1. Considering the case of an optimal strategy for player 1 in the T -stage game $\Gamma_T(p)$, we have shown:

Proposition 3 For p in $\Delta(K)$ and $T \geq 1$,

$$v_T(p) \leq \text{cavu}(p) + \frac{M \sum_{k \in K} \sqrt{p^k(1-p^k)}}{\sqrt{T}}.$$

It remains to conclude about the existence of the uniform value. We have seen that player 1 can guarantee $\text{cavu}(p)$ and that player 2 can guarantee $\lim_T v_T(p)$, and we obtain from Proposition 3 that $\lim_T v_T(p) \leq \text{cavu}(p)$. This is enough to deduce Aumann and Maschler's celebrated "cavu" theorem.

Theorem 1 Aumann and Maschler (1995). The game $\Gamma(p)$ has a uniform value which is $\text{cavu}(p)$.

T -stage Values and the Recursive Formula

As the T -stage game is a zero-sum game with incomplete information where player 1 is informed, we can write:

$$\begin{aligned} v_T(p) &= \inf_{\sigma^2 \in \Sigma^2} \sup_{\sigma^1 \in \Sigma^1} \gamma_T^{1,p}(\sigma), \\ &= \inf_{\sigma^2 \in \Sigma^2} \sup_{\sigma^1 \in \Sigma^1} \sum_{k \in K} p^k \gamma_T^{1, \delta_k}(\sigma), \\ &= \inf_{\sigma^2 \in \Sigma^2} \sum_{k \in K} p^k \left(\sup_{\sigma^1 \in \Sigma^1} \gamma_T^{1, \delta_k}(\sigma) \right). \end{aligned}$$

This shows that v_T is the infimum of a family of affine functions of p , hence is a concave function

of p . This concavity represents the advantage of player 1 to possess the information on the selected state. Clearly, we have $v_T(p) \geq u(p)$; hence, we

get the inequalities: $\forall T \geq 1$, $\text{cavu}(p) \leq v_T(p) \leq \text{cavu}(p) + \frac{M \sum_{k \in K} \sqrt{p^k(1-p^k)}}{\sqrt{T}}$.

It is also easy to prove that the T -stage value functions satisfy the following recursive formula:

$$\begin{aligned} v_{T+1}(p) &= \frac{1}{T+1} \max_{x \in \Delta(I)^K} \min_{y \in \Delta(J)^K} \left(G(p, x, y) + T \sum_{i \in I} x(p)(i) v_T(\hat{p}(x, i)) \right), \\ &= \frac{1}{T+1} \min_{y \in \Delta(J)^K} \max_{x \in \Delta(I)^K} \left(G(p, x, y) + T \sum_{i \in I} x(p)(i) v_T(\hat{p}(x, i)) \right), \end{aligned}$$

where $x = (x^k(i))_{i \in I, k \in K}$, with x^k the mixed action used at stage 1 by player 1 if the state is k , $G(p, x, y) = \sum_{k,i,j} p^k G^k(x^k(i), y(j))$ is the expected payoff of stage 1, $x(p)(i) = \sum_{k \in K} p^k x^k(i)$ is the probability that action i is played at stage 1, and $\hat{p}(x, i)$ is the conditional probability on K given i .

The next property interprets easily: the advantage of the informed player can only decrease as the number of stages increases (for a proof, one can show that $v_{T+1} \leq v_T$ by induction on T , using the concavity of v_T).

Lemma 6 The T -stage value $v_T(p)$ is non-increasing in T .

Vector Payoffs and Approachability

The following model has been introduced by D. Blackwell (1956) and is, strictly speaking, not part of the general definition given in section "Definition of the Subject and Its Importance." We still have a family of $I \times J$ matrices $(G^k)_{k \in K}$, where K is a finite set of parameters. At each stage t , simultaneously player 1 chooses $i_t \in I$ and player 2

chooses $j_t \in J$, and the stage “payoff” is the full vector $G(i_t, j_t) = (G^k(i_t, j_t))_{k \in K}$ in $\mathbb{R}^K$. Notice that there is no initial probability or true state of nature here, and both players have a symmetric role. We assume here that after each stage both players observe exactly the stage vector payoff (but one can check that assuming that the action profiles are observed would not change the results). A natural question is then to determine the sets C in $\mathbb{R}^K$ such that player 1 (for example) can force the average long term payoff to belong to C ? Such sets will be called *approachable* by player 1.

In section “[Vector Payoffs and Approachability](#),” we use Euclidean distances and norms. Denote by $F = \{(G^k(i, j))_{k \in K}, i \in I, j \in J\}$ the finite set of possible stage payoffs and by M a constant such that $\|u\| \leq M$ for each u in F . A strategy for player 1, resp. player 2, is an element $\sigma = (\sigma_t)_{t \geq 1}$, where σ_t maps F^{t-1} into $\Delta(I)$, resp. $\Delta(J)$. Strategy spaces for player 1 and 2 are, respectively, denoted by Σ and $\mathcal{T}$. A strategy profile (σ, τ) naturally induces a unique probability on $(I \times J \times F)^\infty$ denoted by $\mathbb{P}_{\sigma, \tau}$. Let C be a “target” set that will always be assumed, without loss of generality, a closed subset of $\mathbb{R}^K$. We denote by g_t the random variable, with value in F , of the payoff of stage t , and we use $\bar{g}_t = \frac{1}{t} \sum_{t'=1}^t g_{t'} \in \text{conv}(F)$, and finally $d_t = d(\bar{g}_t, C)$ for the distance from $\bar{g}_t$ to C .

Definition 3 C is approachable by player 1 if: $\forall \varepsilon > 0, \exists \sigma \in \Sigma, \exists T, \forall \tau \in \mathcal{T}, \forall t \geq T, \mathbb{E}_{\sigma, \tau}(d_t) \leq \varepsilon$. C is excludable by player 1 if there exist $\delta > 0$ such that $\{z \in \mathbb{R}^K, d(z, C) \geq \delta\}$ is approachable by player 1.

Approachability and excludability for player 2 are defined similarly. C is approachable by player 1 if for each $\varepsilon > 0$, this player can force that for t large we have $\mathbb{E}_{\sigma, \tau}(d_t) \leq \varepsilon$, so the average payoff will be ε -close to C with high probability. A set cannot be approachable by a player as well as excludable by the other player. In the usual case where K is a singleton, we are in dimension 1 and the Minmax theorem implies that for each t , the interval $[t, +\infty]$ is either approachable by player 1 or excludable by player 2, depending on the comparison between t and the value $\max_{x \in \Delta(I)} \min_{y \in \Delta(J)} G(x, y) = \min_{y \in \Delta(J)} \max_{x \in \Delta(I)} G(x, y)$.

Necessary and Sufficient Conditions for Approachability

Given a mixed action x in $\Delta(I)$, we write xG for the set of possible vector payoffs when player 1 uses x , i.e., $xG = \{G(x, y), y \in \Delta(J)\} = \text{conv} \{\sum_{i \in I} x_i G(i, j), j \in J\}$. Similarly, we write $Gy = \{G(x, y), x \in \Delta(I)\}$ for y in $\Delta(J)$.

Definition 4 The set C is a *B(lackwell)-set* for player 1 if for every $z \notin C$, there exists $z' \in C$ and $x \in \Delta(I)$ such that: (i) $\|z' - z\| = d(z, C)$ and (ii) the hyperplane containing z' and orthogonal to $[z, z']$ separates z from xG (Fig. 3).


For example, any set xG , with x in $\Delta(I)$, is a *B-set* for player 1. Given a *B-set* for player 1, we now construct a strategy σ adapted to C as follows. At each positive stage $t + 1$, player 1 considers the current average payoff $\bar{g}_t$. If $\bar{g}_t \in C$, or if $t = 0$, σ plays arbitrarily at stage $t + 1$. Otherwise, σ plays at stage $t + 1$ a mixed action x satisfying the previous definition for $z = \bar{g}_t$.

Theorem 2 If C is a *B-set* for player 1, a strategy σ adapted to C satisfies:

$$\forall \tau \in \mathcal{T}, \forall t \geq 1 \quad \mathbb{E}_{\sigma, \tau}(d_t) \leq \frac{2M}{\sqrt{t}} \quad \text{and} \quad d_t \rightarrow_{t \rightarrow \infty} 0$$

$$\mathbb{P}_{\sigma, \tau} a.s.$$

As an illustration, in dimension 1 and for $C = \{0\}$, this theorem implies that a bounded sequence $(x_t)_{t \geq 1}$ of reals, such that the product



Repeated Games with Incomplete Information, Fig. 3 The Blackwell property

$x_{T+1} \left(\frac{1}{T} \sum_{t=1}^T x_t \right)$ is nonpositive for each T , Cesaro converges to zero.

Proof Assume that player 1 plays σ adapted to C , whereas player 2 plays some strategy τ . Fix $t \geq 1$, and assume that $\bar{g}_t \notin C$. Consider $z' \in C$ and $x \in \Delta(I)$ satisfying (i) and (ii) of Definition 4 for $z = \bar{g}_t$. We have:

$$\begin{aligned} d_{t+1}^2 &= d(\bar{g}_{t+1}, C)^2 \leq \|\bar{g}_{t+1} - z'\|^2 \\ &= \left\| \frac{1}{t+1} \sum_{l=1}^{t+1} g_l - z' \right\|^2 \\ &= \left\| \frac{1}{t+1} (g_{t+1} - z') + \frac{t}{t+1} (\bar{g}_t - z') \right\|^2 \\ &= \left(\frac{1}{t+1} \right)^2 \|g_{t+1} - z'\|^2 + \left(\frac{t}{t+1} \right)^2 d_t^2 \\ &\quad + \frac{2t}{(t+1)^2} \langle g_{t+1} - z', \bar{g}_t - z' \rangle. \end{aligned}$$

By hypothesis, the expectation, given the first t action profiles $h_t \in (I \times J)^t$, of the above scalar product is nonpositive, so $\mathbb{E}(d_{t+1}^2 | h_t) \leq \left(\frac{t}{t+1} \right)^2 d_t^2 + \left(\frac{1}{t+1} \right)^2 \mathbb{E}(\|g_{t+1} - z'\|^2 | h_t)$. Since $\mathbb{E}(\|g_{t+1} - z'\|^2 | h_t) \leq \mathbb{E}(\|g_{t+1} - \bar{g}_t\|^2 | h_t) \leq (2M)^2$, we have:

$$\begin{aligned} \mathbb{E}(d_{t+1}^2 | h_t) &\leq \left(\frac{t}{t+1} \right)^2 d_t^2 \\ &\quad + \left(\frac{1}{t+1} \right)^2 4M^2. \end{aligned} \quad (2)$$

Taking the expectation, we get, whether $\bar{g}_t \notin C$ or not: $\forall t \geq 1, \mathbb{E}(d_{t+1}^2) \leq \left(\frac{t}{t+1} \right)^2 \mathbb{E}(d_t^2) + \left(\frac{1}{t+1} \right)^2 4M^2$. By induction, we obtain that for each $t \geq 1$, $\mathbb{E}(d_t^2) \leq \frac{4M^2}{t}$, and $\mathbb{E}(d_t) \leq \frac{2M}{\sqrt{t}}$.

Put now, as in Sorin (2002), $e_t = d_t^2 + \sum_{t' > t} \frac{4M^2}{t'^2}$. Inequality (2) gives $\mathbb{E}(e_{t+1} | h_t) \leq e_t$, so (e_t) is a nonnegative supermartingale which expectation goes to zero. By a standard probability result, we obtain $e_t \rightarrow_{t \rightarrow \infty} 0$ $\mathbb{P}_{\sigma, \tau}$ a.s., and finally $d_t \rightarrow_{t \rightarrow \infty} 0$ $\mathbb{P}_{\sigma, \tau}$ a.s. $\square$

This theorem implies that any B -set for player 1 is approachable by this player. The converse is true for convex sets.

Theorem 3 Let C be a closed convex subset of $\mathbb{R}^K$.

(i) C is a B -set for player 1,

$\Leftrightarrow$ (ii) $\forall y \in \Delta(J), G_y \cap C \neq \emptyset$,

$\Leftrightarrow$ (iii) C is approachable by player 1,

$\Leftrightarrow$ (iv) $\forall q \in \mathbb{R}^K, \max_{x \in \Delta(I)} \min_{y \in \Delta(J)} \sum_{k \in K} q^k G^k(x, y) \geq$

$\inf_{c \in C} \langle q, c \rangle$.

Proof The implication (i) $\Rightarrow$ (iii) comes from Theorem 2. Proof of (iii) $\Rightarrow$ (ii): assume there exists $y \in \Delta(J)$ such that $G_y \cap C = \emptyset$. Since G_y is approachable by player 2, then C is excludable by player 2 and thus C is not approachable by player 1. Proof of (ii) $\Rightarrow$ (i): Assume that $G_y \cap C \neq \emptyset \forall y \in \Delta(J)$. Consider $z' \notin C$ and define z' as its projection onto C . Define the matrix game where payoffs are projected towards the direction $z' - z$, i.e., the matrix game $\sum_{k \in K} (z'^k - z^k) G^k$. By assumption, one has: $\forall y \in \Delta(J), \exists x \in \Delta(I)$ such that $G(x, y) \in C$, hence such that:

$$\langle z' - z, G(x, y) \rangle \geq \min_{c \in C} \langle z' - z, c \rangle = \langle z' - z, z' \rangle.$$

So $\min_{y \in \Delta(J)} \max_{x \in \Delta(I)} \langle z' - z, G(x, y) \rangle \geq \langle z' - z, z' \rangle$. By the minmax theorem, there exists x in $\Delta(I)$ such that $\forall y \in \Delta(J), \langle z' - z, G(x, y) \rangle \geq \langle z' - z, z' \rangle$, that is $\langle z' - z, z' - G(x, y) \rangle \leq 0$.

(iv) means that any half-space containing C is approachable by player 1.

(iii) $\Rightarrow$ (iv) is thus clear. (iv) $\Rightarrow$ (i) is similar to (ii) $\Rightarrow$ (i). $\square$

Up to minor formulation differences, Theorems 2 and 3 are due to Blackwell (1956). Later on, X. Spinat (2002) proved the following characterization.

Theorem 4 A closed set is approachable for player 1 if and only if it contains a B -set for player 1.

As a consequence, it shows that adding the condition $d_t \rightarrow_{t \rightarrow \infty} 0$ $\mathbb{P}_{\sigma, \tau}$ a.s. in the definition of approachability does not modify the notion.

Approachability for Player 1 Versus Excludability for Player 2

As a corollary of Theorem 3, we obtain that: *A closed convex set in $\mathbb{R}^K$ is either approachable by player 1, or excludable by player 2.*

One can show that when K is a singleton, then any set is either approachable by player 1, or excludable by player 2. A simple example of a set which is neither approachable for player 1 nor excludable by player 2 is given in dimension 2 by:

$$G = \begin{pmatrix} (0,0) & (0,0) \\ (1,0) & (1,1) \end{pmatrix}, \text{ and } C = \{(1/2, v), 0 \leq v \leq 1/4\} \cup \{(1, v), 1/4 \leq v \leq 1\} \text{ (see Sorin 2002)}.$$

Weak Approachability

One can weaken the definition of approachability by giving up time uniformity.

Definition 5 C is weakly approachable by player 1 if: $\forall \varepsilon > 0, \exists T, \forall t \geq T, \exists \sigma \in \Sigma, \forall \tau \in \mathcal{T}, \mathbb{E}_{\sigma, \tau}(d_t) \leq \varepsilon$. C is weakly excludable by player 1 if there exists $\delta > 0$ such that $\{z \in \mathbb{R}^K, d(z, C) \geq \delta\}$ is weakly approachable by player 1.

N. Vieille (1992) has proved, via the consideration of certain differential games:

Theorem 5 A subset of $\mathbb{R}^K$ is either weakly approachable by player 1 or weakly excludable by player 2.

Back to the Standard Model

Let us come back to Aumann and Maschler's model with a finite family of matrices $(G^k)_{k \in K}$ and an initial probability p on $\Delta(K)$. By Theorem 1, the repeated game $\Gamma(p)$ has a uniform value which is $\text{cavu}(p)$, and Blackwell approachability will allow for the construction of an *explicit* optimal strategy for the uninformed player. Considering a hyperplane which is tangent to cavu at p , we can find a vector l in $\mathbb{R}^K$ such that

$$\langle l, p \rangle = \text{cavu}(p) \quad \text{and} \quad \forall q \in \Delta(K), \langle l, q \rangle \geq \text{cavu}(q) \geq u(q).$$

Define now the orthant $C = \{z \in \mathbb{R}^K, z^k \leq l^k \forall k \in K\}$. Recall that player 2 does not know the selected state, and an optimal strategy for him cannot depend on player 1's strategy and consequently on a martingale *a posteriori*. He will play in a way such that player 1's long term payoff is, *simultaneously for each k in K* , not greater than l^k if the state is k .

Fix $q = (q^k)_{k \in K}$ in $\mathbb{R}^K$. If there exists k with $q^k > 0$, we clearly have $\inf_{c \in C} \langle q, c \rangle < \langle q, p \rangle = -\infty \leq \max_{y \in \Delta(J)} \min_{x \in \Delta(I)} \sum_{k \in K} q^k G^k(x, y)$. Assume now that $q^k \leq 0$ for each k , with $q \neq 0$. Write $s = \sum_{k \in K} (-q^k)$.

$$\begin{aligned} \inf_{c \in C} \langle q, c \rangle &= \sum_{k \in K} q^k l^k \\ &= -s \langle l, \frac{-q}{s} \rangle \\ &\leq -s u\left(\frac{-q}{s}\right) \\ &\leq -s \max_{x \in \Delta(I)} \min_{y \in \Delta(J)} \sum_{k \in K} \frac{-q^k}{s} G^k(x, y) \\ &= \max_{y \in \Delta(J)} \min_{x \in \Delta(I)} \sum_{k \in K} q^k G^k(x, y) \end{aligned}$$

This is condition (iv) of Theorem 3, adapted to player 2. So C is a B -set for player 2, and a strategy τ adapted to C satisfies by Theorem 2: $\forall \sigma \in \Sigma, \forall k \in K$,

$$\begin{aligned} \mathbb{E}_{\sigma, \tau} \left(\frac{1}{T} \sum_{t=1}^T G^k(\tilde{i}_t, \tilde{j}_t) - l^k \right) \\ \leq \mathbb{E}_{\sigma, \tau} \left(d \left(\frac{1}{T} \sum_{t=1}^T G^k(\tilde{i}_t, \tilde{j}_t), C \right) \right) \leq \frac{2M}{\sqrt{T}}, \end{aligned}$$

(where M is here an upper bound for the Euclidean norms of the vectors, $(G^k(i; j))_{k \in K}$ with $i \in I$ and $j \in J$). So,

$$\begin{aligned} \gamma_T^{1, p}(\sigma, \tau) &= \sum_{k \in K} p^k \left(\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\sigma, \tau} (G^k(\tilde{i}_t, \tilde{j}_t)) \right) \\ &\leq \langle p, l \rangle + \frac{2M}{\sqrt{T}} = \text{cavu}(p) + \frac{2M}{\sqrt{T}}. \end{aligned}$$

As shown by Kohlberg (1975), the approachability strategy τ is thus an optimal strategy for player 2 in the repeated game $\Gamma(p)$.

No-Regret Strategies

The theory of approachability can also be used to prove the existence of "no-regret strategies."

Consider a decision-maker, who has to select at each stage n some action i_n in a finite set I . The environment (nature, adversary, other agents following their own goals) will select a stochastic process $(j_n)_{n \geq 1}$ with values in a finite set J , and

the decision-maker knows a priori nothing about the way the sequence $(j_n)_n$ is chosen. There is a given payoff function $g: I \times J \rightarrow \mathbb{R}$, known by the decision-maker, and at the end of each stage n the decision-maker observes j_n and receives the payoff $g(i_n, j_n)$.

Basic Example $I = J = \{0, 1\}$, and $g(i, j) = 1$ if and only if $i = j$: player 1 tries to guess at each stage n the value j_n (a stage could correspond to a day, $j_n = 0$ meaning there is no rain on day n , $j_n = 1$ meaning there is some rain on day n).

What means to play well for the decision-maker? In the basic example, is it good to guess correctly 90% of the stages? Probably not if it happens that $j_n = 1$ for each n .

A strategy for the decision-maker is an element $\sigma = (\sigma_t)_{t \geq 1}$, where for each t σ_t is a mapping from $(I \times J)^{t-1}$ to $\Delta(I)$. A strategy for Nature is an element $\tau = (\tau_t)_{t \geq 1}$, where for each t τ_t is a mapping from $(I \times J)^{t-1}$ to $\Delta(J)$. The sets of strategies of the decision-maker and Nature are, respectively, denoted by Σ and $\mathcal{T}$, and a strategy profile (σ, τ) naturally induces a unique probability on $(I \times J)^\infty$ denoted by $\mathbb{P}_{\sigma, \tau}$. For each stage n , $g_n = g(i_n, j_n)$ is the random variable of the payoff of stage n , and $\bar{g}_n = \frac{1}{n} \sum_{t=1}^n g_t$.

Suppose that at the end of some stage n , $(i_1, j_1, \dots, i_n, j_n)$ has been played. The average payoff for the decision-maker is $\bar{g}_n$, and he can compare this payoff with the payoff he would have got if he had played constantly some action i in I . The difference $\frac{1}{n} \sum_{t=1}^n g(i, j_t) - \bar{g}_n$ is called the regret of the decision-maker for not having played constantly action i .

Definition 6 A strategy σ of the decision-maker has no external regret if for all strategy τ of Nature,

$$\limsup_{n \rightarrow \infty} \left(\max_{i \in I} \frac{1}{n} \sum_{t=1}^n g(i, j_t) - \bar{g}_n \right) \leq 0 \quad \mathbb{P}_{\sigma, \tau} a.s.$$

We now define the stronger notion of internal regret: for each pair of actions i and l in I , we do not want the decision-maker to regret to have

played action l at each stage where he actually played action i . For $n \geq 1$, i and l in I , let us introduce the random variable:

$$\bar{R}_n(i, l) = \frac{1}{n} \sum_{t \in \{1, \dots, n\}, i_t = i} (g(l, j_t) - g(i, j_t)).$$

Definition 7 A strategy σ of the decision-maker has no internal regret if for each strategy τ of Nature,

$$\max_{i \in I, l \in I} \bar{R}_n(i, l) \xrightarrow{n \rightarrow \infty} 0 \quad \mathbb{P}_{\sigma, \tau} a.s.$$

Theorem 6 There exists a strategy σ of the decision-maker with no internal regret.

Proof Define $K = I \times I$. Consider the dynamic game with vector payoffs where at each stage n player 1 chooses i_n in I , player 2 chooses j_n in J , and the vector payoff in $\mathbb{R}^K$ is $(r(i_n, j_n))_{i, l}$ with:

$$\forall (i, l) \in K, \quad r(i_n, j_n)_{i, l} = \begin{cases} g(l, j_n) - g(i, j_n) & \text{if } i_n = i \\ 0 & \text{if } i_n \neq i \end{cases}$$

If $i_n = i$, $r(i_n, j_n)_{i, l}$ is the difference between the payoff that the decision-maker could have got at stage n by playing l and what he actually got. We denote by $r_n = r(i_n, j_n)$ the vector of regrets at stage n in $\mathbb{R}^K$, and we write $\bar{r}_n = \frac{1}{n} \sum_{t=1}^n r_t$. Notice that $\bar{r}_n$ is nothing but the regret vector $\bar{R}_n$.

Let $C = \mathbb{R}_-^K$ be the negative orthant of $\mathbb{R}^K$. We use the Euclidean norm in this proof. For each r in $\mathbb{R}^K$, the projection of r to C is given by $\pi_C(r) = (\min\{r(k), 0\})_{k \in K}$, and $d(r, C) = \|r^+\|$, where $r^+ = (\max\{r(k), 0\})_{k \in K}$.

We now show that condition (ii) of Theorem 3 is satisfied. Fix $y = (y_j)_{j \in J}$ in $\Delta(J)$ and consider i^* achieving $\max_{i \in I} g(i, y)$. $r(i^*, y)_{i, l} = 0$ if $i \neq i^*$, and $r(i^*, y)_{i^*, l} = g(l, y) - g(i^*, y) \leq 0$. So C is a B-set by Theorem 3, and by Theorem 2 we have the existence of a strategy σ of the decision-maker such that for each strategy τ of Nature:

$$\forall n \geq 1, \mathbb{E}_{\sigma, \tau}(\|\bar{r}_n^+\|) \leq \frac{2M}{\sqrt{n}} \quad \text{and} \\ \|\bar{r}_n^+\| \xrightarrow{n \rightarrow \infty} 0 \quad \mathbb{P}_{\sigma, \tau} \text{ a.s.}$$

where M is the constant $\max_{i,j} \|r(i, j)\|$. This concludes the proof of Theorem 6: for each strategy τ of Nature, stage n and all i and l in I ,

$$\mathbb{E}_{\sigma, \tau} \bar{R}_n(i, l) \leq \frac{2M}{\sqrt{n}} \quad \text{and} \\ \max_{i \in I, l \in I} \bar{R}_n(i, l) \xrightarrow{n \rightarrow \infty} 0 \quad \mathbb{P}_{\sigma, \tau}.$$

Zero-Sum Games with Lack of Information on Both Sides

The following model has also been introduced by Aumann and Maschler (1995). We are still in the context of zero-sum repeated games with observable actions, but it is no longer assumed that one of the players is fully informed. The set of states is here a product $K \times L$ of finite sets, and we have a family of matrices $(G^{k,l})_{(k,l) \in K \times L}$ with size $I \times J$, as well as initial probabilities p on K , and q on L . In the game $\Gamma(p, q)$, a state of nature (k, l) is first selected according to the product probability $p \otimes q$, then k , resp. l , is announced to player 1, resp. player 2 only. Then the matrix game $G^{k,l}$ is repeated over and over: at every stage, simultaneously player 1 chooses a row i in I , whereas player 2 chooses a column j in J , the stage payoff for player 1 is $G^{k,l}(i, j)$, but only i and j are publicly announced before proceeding to the next stage.

The average payoff for player 1 in the T -stage game is written: $\gamma_T^{1,p,q}(\sigma^1, \sigma^2) = \mathbb{E}_{\sigma^1, \sigma^2}^{p,q} \left(\frac{1}{T} \sum_{t=1}^T G^{\tilde{i}_t, \tilde{j}_t}(\tilde{i}_t, \tilde{j}_t) \right)$, and the T -stage value is written $v_T(p, q)$. Similarly, the λ -discounted value of the game will be written $v_\lambda(p, q)$.

The nonrevealing game now corresponds to the case where player 1 plays independently of k and player 2 plays independently of l . Its value is denoted by:

$$u(p, q) = \max_{x \in \Delta(I)} \min_{y \in \Delta(J)} \sum_{k,l} p^k q^l G^{k,l}(x, y). \quad (3)$$

Given a continuous function $f: \Delta(K) \times \Delta(L) \rightarrow \mathbb{R}$, we denote by $\text{cav}_I f$ the concavification of f with

respect to the first variable: for each (p, q) in $\Delta(K) \times \Delta(L)$, $\text{cav}_I f(p, q)$ is the value at p of the smallest concave function from $\Delta(K)$ to $\mathbb{R}$ which is above $f(\cdot, q)$. Similarly, we denote by $\text{vex}_I f$ the convexification of f with respect to the second variable. It can be shown that $\text{cav}_I f$ and $\text{vex}_I f$ are continuous, and we can compose $\text{cav}_I \text{vex}_I f$ and $\text{vex}_I \text{cav}_I f$. These functions are both concave in the first variable and convex in the second variable, and they satisfy $\text{cav}_I \text{vex}_I f(p, q) \leq \text{vex}_I \text{cav}_I f(p, q)$.

Maxmin and Minmax of the Repeated Game

Theorem 1 generalizes as follows.

Theorem 7 Aumann and Maschler (1995) In the repeated game $\Gamma(p, q)$, the greatest quantity which can be guaranteed by player 1 is $\text{cav}_I \text{vex}_I u(p, q)$, and the smallest quantity which can be guaranteed by player 2 is $\text{vex}_I \text{cav}_I u(p, q)$.

Aumann, Maschler, and Stearns also showed that $\text{cav}_I \text{vex}_I u(p, q)$ can be *defended* by player 2, uniformly in time, i.e., that $\forall \varepsilon > 0, \forall \sigma^1, \exists T_0, \exists \sigma^2, \forall T \geq T_0, \gamma_T^{p,q}(\sigma^1, \sigma^2) \leq \text{cav}_I \text{vex}_I u(p, q) + \varepsilon$. Similarly, $\text{vex}_I \text{cav}_I u(p, q)$ can be *defended* by player 1.

The proof uses the martingales of *a posteriori* of each player, and a useful notion is that of the informational content of a strategy: for a strategy σ^1 of the first player, it is defined as: $I(\sigma^1) = \sup_{\sigma^2} \mathbb{E}_{\sigma^1, \sigma^2}^{p,q} \left(\sum_{k \in K} \sum_{t=0}^{\infty} (p_{t+1}^k(\sigma^1) - p_t^k(\sigma^1))^2 \right)$, where $p_t(\sigma^1)$ is the *a posteriori* on K of player 2 after stage t given that player 1 uses σ^1 . By linearity of the expectation, the supremum can be restricted to strategies of player 2 which are both pure and independent of l .

Theorem 7 implies that $\text{cav}_I \text{vex}_I u(p, q) = \sup_{\sigma^1 \in \Sigma^1} \liminf_T \left(\inf_{\sigma^2 \in \Sigma^2} \gamma_T^{1,p,q}(\sigma^1, \sigma^2) \right)$, and $\text{cav}_I \text{vex}_I u(p, q)$ is called the *maxmin* of the repeated game $\Gamma(p, q)$. Similarly, $\text{vex}_I \text{cav}_I u(p, q) = \inf_{\sigma^2 \in \Sigma^2} \limsup_T \left(\sup_{\sigma^1 \in \Sigma^1} \gamma_T^1(\sigma^1, \sigma^2) \right)$ is called the *minmax* of $\Gamma(p, q)$. As a corollary, we obtain that the repeated game $\Gamma(p, q)$ has a uniform value if and only if: $\text{cav}_I \text{vex}_I u(p, q) = \text{vex}_I \text{cav}_I u(p, q)$. This is not always the case, and there

exist counter-examples to the existence of the uniform value.

Example 4 $K = \{a, a'\}$, and $L = \{b, b'\}$, with p and q uniform.

$$\begin{aligned} G^{a,b} &= \begin{pmatrix} 0 & 0 & 0 & 0 \\ -1 & 1 & 1 & -1 \end{pmatrix} \\ G^{a,b'} &= \begin{pmatrix} 1 & -1 & 1 & -1 \\ 0 & 0 & 0 & 0 \end{pmatrix} \\ G^{a',b} &= \begin{pmatrix} -1 & 1 & -1 & 1 \\ 0 & 0 & 0 & 0 \end{pmatrix} \\ G^{a',b'} &= \begin{pmatrix} 0 & 0 & 0 & 0 \\ 1 & -1 & -1 & 1 \end{pmatrix} \end{aligned}$$

Mertens and Zamir (1971) have shown that here, $\text{cav}_I \text{vex}_{II} u(p, q) = -\frac{1}{4} < 0 = \text{vex}_{II} \text{cav}_I u(p, q)$.

Limit Values

It is easy to see that for each T and λ , the value functions v_T and v_λ are concave in the first variable and convex in the second variable. They are all Lipschitz functions, with the same constant $M = \max_{i,j,k,l} |G^{k,l}(i, j)|$, and here also, recursive formulae can be given. In the following result, v_T and v_λ are viewed as elements of the set $\mathcal{C}$ of continuous mappings from $\Delta(K) \times \Delta(L)$ to $\mathbb{R}$.

Theorem 8 Mertens and Zamir (1971) $(v_T)_T$, as T goes to infinity, and $(v_\lambda)_\lambda$, as λ goes to zero, both uniformly converge to the unique solution f of the following system:

$$\begin{cases} f = \text{vex}_{II} \max\{u, f\} \\ f = \text{cav}_I \min\{u, f\} \end{cases}$$

And the convergence of (v_T) , resp. (v_λ) is in $O(1/\sqrt{T})$, resp. $O(\lambda)$.

The above system can be fruitfully studied without reference to repeated games (see Laraki 2001a, b; Mertens and Zamir 1977; Sorin 1984b).

Remark Let U be the set of all nonrevealing value functions, i.e., of functions from $\Delta(K) \times \Delta(L)$ to $\mathbb{R}$ satisfying Eq. (3) for some family of matrices

$(G^{k,l})_{k,l}$. One can easily show that any mapping in $\mathcal{C}$ is a uniform limit of elements in U .

Correlated Initial Information

A more general model can be written, where it is no longer assumed that the initial information of the players is independent. The set of states is now denoted by R (instead of $K \times L$), initially a state r in R is chosen according to a known probability $p = (p^r)_{r \in R}$, and each player receives a deterministic signal depending on r . Equivalently, each player i has a partition R^i of R and observes the element of his partition which contains the selected state.

After the first stage, player 1 will play an action $x = (x^r)_{r \in R}$ which is measurable with respect to R^1 , i.e., $(r \rightarrow x^r)$ is constant on each atom of R^1 . After having observed player 1's action at the first stage, the conditional probability on R necessarily belongs to the set:

$$\Pi^1(p) = \left\{ (\alpha^r p^r)_{r \in R}, \forall r \alpha^r \geq 0, \sum_r \alpha^r p^r = 1 \text{ and } (\alpha^r)_r \text{ is } R^1\text{-measurable} \right\}.$$

$\Pi^1(p)$ contains p and is a convex compact subset of $\Delta(R)$. A mapping f from $\Delta(R)$ to $\mathbb{R}$ is now said to be I-concave if for each p in $\Delta(R)$, the restriction of f to $\Pi^1(p)$ is concave. And given $g: \Delta(R) \rightarrow \mathbb{R}$ which is bounded from above, we define the concavification $\text{cav}_I g$ as the smallest function above g which is I-concave. Similarly one can define the set $\Pi^{II}(p)$ and the notions of II-convexity and II-convexification. With these generalized definitions, the results of Theorem 7 and 8 perfectly extend (Mertens and Zamir 1971).

Nonzero-sum Games with Lack of Information on One Side

We now consider the generalization of the standard model of section “The Standard Model of Aumann and Maschler” to the nonzero-sum case. Hence, two players infinitely repeat the same bimatrix game, with player 1 only knowing the

bimatrix. Formally, we have a finite set of states K , an initial probability p on K , and families of $I \times J$ -payoff matrices $(A^k)_{k \in K}$ and $(B^k)_{k \in K}$. Initially, a state k in K is selected according to p , and announced to player 1 only. Then the bimatrix game (A^k, B^k) is repeated over and over: at every stage, simultaneously player 1 chooses a row i in I , whereas player 2 chooses a column j in J , the stage payoff for player 1 is then $A^k(i, j)$, the stage payoff for player 2 is $B^k(i, j)$, but only i and j are publicly announced before proceeding to the next stage. Without loss of generality, we assume that $p^k > 0$ for each k and that each player has at least 2 actions.

Given a strategy pair (σ^1, σ^2) , it is here convenient to denote the expected payoffs up to stage T by:

$$\begin{aligned} \alpha_T^p(\sigma^1, \sigma^2) &= \mathbb{E}_{p, \sigma^1, \sigma^2} \left(\frac{1}{T} \sum_{t=1}^T A^{\tilde{k}}(\tilde{i}_t, \tilde{j}_t) \right) \\ &= \sum_{k \in K} p^k \alpha_T^k(\sigma^1, \sigma^2). \end{aligned}$$

$$\begin{aligned} \beta_T^p(\sigma^1, \sigma^2) &= \mathbb{E}_{p, \sigma^1, \sigma^2} \left(\frac{1}{T} \sum_{t=1}^T B^{\tilde{k}}(\tilde{i}_t, \tilde{j}_t) \right) \\ &= \sum_{k \in K} p^k \beta_T^k(\sigma^1, \sigma^2). \end{aligned}$$

Given a probability q on K , we write $A(q) = \sum_k q^k A^k$, $B(q) = \sum_k q^k B^k$, $u(q) = \max_{x \in \Delta(I)} \min_{y \in \Delta(J)} A(q)(x, y)$ and $v(q) = \max_{y \in \Delta(J)} \min_{x \in \Delta(I)} B(q)(x, y)$. If $\gamma = (\gamma(i, j))_{(i, j) \in I \times J} \in \Delta(I \times J)$, we put $A(q)(\gamma) = \sum_{(i, j) \in I \times J} \gamma(i, j) A(q)(i, j)$ and similarly $B(q)(\gamma) = \sum_{(i, j) \in I \times J} \gamma(i, j) B(q)(i, j)$.

Existence of Equilibria

The question of existence of an equilibrium has remained unsolved for long. Sorin (1983) proved the existence of an equilibrium for two states of nature, and the general case has been solved by Simon et al. (1995).

Exactly as in the zero-sum case, a strategy pair σ induces a sequence of *a posteriori* $(p_t(\sigma))_{t \geq 0}$ which is a $\mathbb{P}_{p, \sigma}$ -martingale with values in $\Delta(K)$. We will concentrate on the cases where this martingale moves only once.

Definition 8 A joint plan is a triple (S, λ, γ) , where:

- S is a finite non empty set (of messages),
- $\lambda = (\lambda^k)_{k \in K}$ (signaling strategy) with for each k , $\lambda^k \in \Delta(S)$ and for each s , $\lambda_s =_{\text{def}} \sum_{k \in K} p^k \lambda_s^k > 0$,
- $\gamma = (\gamma_s)_{s \in S}$ (contract) with for each s , $\gamma_s \in \Delta(I \times J)$.

The idea is due to Aumann, Maschler, and Stearns. Player 1 observes k , then chooses $s \in S$ according to λ^k and announces s to player 2. Then the players play pure actions corresponding to the frequencies $\gamma_s(i, j)$, for i in I and j in J . Given a joint plan (S, λ, γ) , we define:

- $\forall s \in S$, $p_s = (p_s^k)_{k \in K} \in \Delta(K)$, with $p_s^k = \frac{p^k \lambda_s^k}{\lambda_s}$ for each k . p_s is the *a posteriori* on K given s .
- $\varphi = (\varphi^k)_{k \in K} \in \mathbb{R}^K$, with for each k , $\varphi^k = \max_{s \in S} A^k(\gamma_s)$.
- $\forall s \in S$, $\psi_s = B(p_s)(\gamma_s)$ and $\psi = \sum_{k \in K} p^k \sum_{s \in S} \lambda_s^k B^k(\gamma_s) = \sum_{s \in S} \lambda_s \psi_s$.

Definition 9 A joint plan (S, λ, γ) is an equilibrium joint plan if:

- (i) $\forall s \in S$, $\psi_s \geq \text{vexv}(p_s)$
- (ii) $\forall k \in K$, $\forall s \in S$ s.t. $p_s^k > 0$, $A^k(\gamma_s) = \varphi^k$
- (iii) $\forall q \in \Delta(K)$, $\langle \varphi, q \rangle \geq u(q)$

Condition (ii) can be seen as an incentive condition for player 1 to choose s according to λ^k . Given an equilibrium joint plan (S, λ, γ) , one define a strategy pair $(\sigma^{1*}, \sigma^{2*})$ adapted to it. For each message s , first fix a sequence $(i_t^s, j_t^s)_{t \geq 1}$ of elements in $I \times J$ such that for each (i, j) , the empirical frequencies converge to the corresponding probability: $\frac{1}{T} |\{t, 1 \leq t \leq T, (i_t^s, j_t^s) = (i, j)\}| \rightarrow_{T \rightarrow \infty} \gamma_s(i, j)$. We also fix an injective mapping f from S to I^l , where l is large enough, corresponding to a code between the players to announce an element in S . σ^{1*} is precisely defined as follows. Player 1 observes the selected state k , then chooses s according to λ^k , and announces s to player 2 by playing $f(s)$ at the first l stages. Finally, σ^{1*} plays i_t^s at each stage $t > l$ as long as player 2 plays j_t^s . If at some stage $t > l$ player 2 does not play j_t^s , then

player 1 punishes his opponent by playing an optimal strategy in the zero-sum game with initial probability p_s and payoffs for player 1 given by $(-B^k)_{k \in K}$. We now define σ^{2*} . Player 2 arbitrarily plays at the beginning of the game, then compute at the end of stage l the message s sent by player 1. Next he plays at each stage $t > l$ the action $\tilde{r}_t^s$ as long as player 1 plays $\tilde{r}_t^s$. If at some stage $t > l$, player 1 does not play $\tilde{r}_t^s$, or if the first l actions of player 1 correspond to no message, then player 2 plays a punishing strategy σ^{-2} such that: $\forall \varepsilon > 0$, $\exists T_0$, $\forall T \geq T_0$, $\forall \sigma^1 \in \Sigma^1$, $\forall k \in K$, $\alpha_k(\sigma^1, \sigma^{-2}) \leq \varphi^k + \varepsilon$. Such a strategy σ^{-2} exists because of condition (iii): it is an approachability strategy for player 2 of the orthant $\{x \in \mathbb{R}^K, \forall k \in K, x^k \leq \varphi^k\}$ (see section “Back to the Standard Model”).

Lemma 7 Sorin (1983) A strategy pair adapted to an equilibrium joint plan is a uniform equilibrium of the repeated game.

Proof The payoffs induced by $(\sigma^{1*}, \sigma^{2*})$ can be easily computed:

$\forall k, \alpha_T^k(\sigma^{1*}, \sigma^{2*}) \rightarrow_{T \rightarrow \infty} \sum_{s \in S} \lambda_s^k A^k(\gamma_s) = \varphi^k$ because of (ii), and $\beta_T^p(\sigma^{1*}, \sigma^{2*}) \rightarrow_{T \rightarrow \infty} \sum_{k \in K} p^k \sum_{s \in S} \lambda_s^k B^k(\gamma_s) = \psi$.

Assume that player 2 plays σ^{2*} . The existence of σ^{2*} implies that no detectable deviation of player 1 is profitable, so if the state is k , player 1 will gain no more than $\max_{s' \in S} A^k(\gamma_{s'})$. But this is just φ^k . The proof can be made uniform in σ^1 and we obtain: $\forall \varepsilon > 0 \exists T_0 \forall T \geq T_0, \forall k \in K, \forall \sigma^1 \in \Sigma^1, \alpha_T^k(\sigma^1, \sigma^{2*}) \leq \varphi^k + \varepsilon$. Finally assume that player 1 plays σ^{1*} . Condition (i) implies that if player 2 uses σ^{2*} , the payoff of this player will be at least $\text{vex } v(p_s)$ if the message is s . Since $\text{vex } v(p_s) (= -\text{cav}(-v(p_s)))$ is the value, from the point of view of player 2 with payoffs $(B^k)_{k \in K}$, of the zero-sum game with initial probability p_s , player 2 fears the punishment by player 1, and $\forall \varepsilon > 0, \exists T_0, \forall T \geq T_0, \forall \sigma^2 \in \Sigma^2, \beta_T^p(\sigma^{1*}, \sigma^2) \leq \sum_{s \in S} \lambda_s \psi_s + \varepsilon = \psi + \varepsilon$. $\square$

To prove the existence of equilibria, we then look for equilibrium joint plans. The first idea is to consider, for each probability r on K , the set of payoff vectors φ compatible with r being an a

posteriori. This leads to the consideration of the following correspondence (for each r , $\Phi(r)$ is a subset of $\mathbb{R}^K$):

$$\Phi : \Delta(K) \rightrightarrows \mathbb{R}^K$$

$r \mapsto \{(A^k(\gamma))_{k \in K}, \text{ where } \gamma \in \Delta(I \times J) \text{ satisfies } B(r)(\gamma) \geq \text{vex } v(r)\}$. It is easy to see that the graph of Φ , i.e., the set $\{(r, \varphi) \in \Delta(K) \times \mathbb{R}^K, \varphi \in \Phi(r)\}$, is compact that Φ has nonempty convex values and satisfies: $\forall r \in \Delta(K), \forall q \in \Delta(K), \exists \varphi \in \Phi(r), < \varphi, q > \geq u(q)$.


Assume now that one can find a finite family $(p_s)_{s \in S}$ of probabilities on K , as well as vectors φ and, for each s , φ_s in $\mathbb{R}^K$ such that: (1) $p \in \text{conv } \{p_s, s \in S\}$, (2) $< \varphi, q > \geq u(q) \forall q \in \Delta(K)$, (3) $\forall s \in S, \varphi_s \in \Phi(p_s)$, and (4) $\forall s \in S, \forall k \in K, \varphi_s^k \leq \varphi^k$ with equality if $p_s^k > 0$. It is then easy to construct an equilibrium joint plan. Thus, we get interested in proving the following result.

Proposition 4 Let p be in $\Delta(K)$, $u: \Delta(K) \rightarrow \mathbb{R}$ be a continuous mapping, and $\Phi: \Delta(K) \rightrightarrows \mathbb{R}^K$ be a correspondence with compact graph and non-empty convex values such that: $\forall r \in \Delta(K), \forall q \in \Delta(K), \exists \varphi \in \Phi(r), < \varphi, q > \geq u(q)$. Then there exists a finite family $(p_s)_{s \in S}$ of elements of $\Delta(K)$, as well as vectors φ and, for each s , φ_s in $\mathbb{R}^K$ such that:

- $p \in \text{conv } \{p_s, s \in S\}$,
- $< \varphi, q > \geq u(q) \forall q \in \Delta(K)$,
- $\forall s \in S, \varphi_s \in \Phi(p_s)$,
- $\forall s \in S, \forall k \in K, \varphi_s^k \leq \varphi^k$ with equality if $p_s^k > 0$.

The proof of Proposition 4 relies, as explained in Renault (2000) or Simon (2002), on a fixed point theorem of Borsuk-Ulam type proved by Simon et al. (1995) via tools from algebraic geometry. A simplified version of this fixed point theorem can be written as follows:

Theorem 9 Simon et al. (1995): Let C be a compact subset of an n -dimensional Euclidean space, $x \in C$ and Y be a finite union of affine subspaces of dimension $n - 1$ of an Euclidean space. Let F be a correspondence from C to Y with



Repeated Games with Incomplete Information, Fig. 4 A Borsuk-Ulam type theorem by Simon, Spiez, and Toruńczyk

compact graph and nonempty convex values. Then there exists $L \subset \partial C$ and $y \in Y$ such that: $\forall l \in L, y \in F(l)$, and $x \in \text{conv}(L)$ (Fig. 4).

Notice that for $n = 1$ (corresponding to 2 states of nature), the image by F of the connected component of C containing x necessarily is a singleton; hence, the result is clear. In the general case, one finally obtains:

Theorem 10 Simon et al. (1995): There exists an equilibrium joint plan. Thus, there exists a uniform equilibrium in the repeated game $\Gamma(p)$.

Characterization of Equilibrium Payoffs

Characterizing equilibrium payoffs, as the Folk theorem does for repeated games with complete information, has been a challenging problem. We denote here by p_0 the initial probability in the interior of $\Delta(K)$. We are interested in the set of equilibrium payoffs, in the convenient following sense:

Definition 10 A vector (a, b) in $\mathbb{R}^K \times \mathbb{R}$ is called an equilibrium payoff of the repeated game $\Gamma(p_0)$ if there exists a strategy pair $(\sigma^{1*}, \sigma^{2*})$ satisfying:

(i) $\forall \varepsilon > 0 \exists T_0 \forall T \geq T_0, \forall k \in K, \forall \sigma^1 \in \Sigma^1, \alpha_T^k(\sigma^1, \sigma^{2*}) \leq \alpha_T^k(\sigma^{1*}, \sigma^{2*}) + \varepsilon, \forall \varepsilon > 0 \exists T_0 \forall T \geq T_0, \forall \sigma^2 \in \Sigma^2, \beta_T^{p_0}(\sigma^{1*}, \sigma^2) \leq \beta_T^{p_0}(\sigma^{1*}, \sigma^{2*}) + \varepsilon$, and (ii) $(\alpha_T^k(\sigma^{1*}, \sigma^{2*}))_{k,T}$ and $(\beta_T^{p_0}(\sigma^{1*}, \sigma^{2*}))_T$ respectively converge to a and b .

Since p lies in the interior of $\Delta(K)$, the first line of (i) is equivalent to: $\forall \varepsilon > 0 \exists T_0 \forall T \geq T_0, \forall \sigma^1 \in \Sigma^1, \alpha_T^p(\sigma^1, \sigma^{2*}) \leq \alpha_T^p(\sigma^{1*}, \sigma^{2*}) + \varepsilon$. The strategy pair

$(\sigma^{1*}, \sigma^{2*})$ is thus a uniform equilibrium of the repeated game, with the additional requirement that expected average payoffs of player 1 converge in each state k . In some sense, player 1 is viewed here as $|K|$ different types or players, and we require the existence of the limit payoff of each type. We will only consider such uniform equilibria in the sequel.

Notice that the above definition implies: $\forall k \in K, \forall \varepsilon > 0, \exists T_0, \forall T \geq T_0, \forall \sigma^1 \in \Sigma^1, \alpha_T^k(\sigma^1, \sigma^{2*}) \leq \alpha_T^k + \varepsilon$. So the orthant $\{x \in \mathbb{R}^K, x^k \leq \alpha_T^k \forall k \in K\}$ is approachable by player 2, and by Theorem 3 and subsection “Back to the Standard Model” one can obtain that:

$$\langle a, q \rangle \geq u(q) \quad \forall q \in \Delta(K) \quad (4)$$

Condition (4) is called the individual rationality condition for player 1 and does not depend on the initial probability in the interior of $\Delta(K)$. Regarding player 2, we have: $\forall \varepsilon > 0 \exists T_0 \forall T \geq T_0, \forall \sigma^2 \in \Sigma^2, \beta_T^{p_0}(\sigma^{1*}, \sigma^2) \leq \beta + \varepsilon$, so by Theorem 1:

$$\beta \geq \text{vex } v(p_0). \quad (5)$$

Condition (5) is the individual rationality condition for player 2: at equilibrium, this player should have at least the value of the game where player 1's plays in order to minimize player 2's payoffs.

Imagine now that σ^{1*} is a nonrevealing strategy for player 1 and that the players play actions with empirical frequencies corresponding to a given probability distribution $\pi = (\pi_{i,j})_{(i,j) \in I \times J} \in \Delta(I \times J)$. We will have: $\forall k \in K, \alpha^k = \sum_{i,j} \pi_{i,j} A^k(i, j)$ and $\beta = \sum_k p_0^k \sum_{i,j} \pi_{i,j} B^k(i, j)$, and if the individual rationality conditions are satisfied, no detectable deviation of a player can be profitable. This leads to the definition of the following set, where M is the constant $\max\{|A^k(i, j)|, |B^k(i, j)|, (i, j) \in I \times J\}$, and $\mathbb{R}_M = [-M, M]$.

Definition 11 Let G be the set of triples $(a, \beta, p) \in \mathbb{R}_M^K \times \mathbb{R}_M \times \Delta(K)$ satisfying:

1. $\forall q \in \Delta(K), \langle a, q \rangle \geq u(q)$,
2. $\beta \geq \text{vex } v(p)$,

3. $\exists \pi \in \Delta(I \times J)$ s.t. $\beta = \sum_k p^k \sum_{i,j} \pi_{i,j} B^k(i, j)$ and $\forall k \in K, a^k \geq \sum_{i,j} \pi_{i,j} A^k(i, j)$ with equality if $p^k > 0$.

We need to considerate every possible initial probability because the main state variable of the model is, here also, the belief, or *a posteriori*, of player 2 on the state of nature. $\{(a, \beta), (a, \beta, p_0) \in G\}$ is the set of payoffs of nonrevealing equilibria of $\Gamma(p_0)$. The importance of the following definition will appear with Theorem 11 below (which unfortunately has not led to a proof of existence of equilibrium payoffs).

Definition 12 G^* is defined as the set of elements $g = (a, \beta, p) \in \mathbb{R}_M^K \times \mathbb{R}_M \times \Delta(K)$ such that there exist a probability space $(\Omega, \mathcal{A}, \mathcal{Q})$, an increasing sequence $(\mathcal{F}_n)_{n \geq 1}$ of finite sub- σ -algebras of $\mathcal{A}$, and a sequence of random variables $(g_n)_{n \geq 1} = (a_n, \beta_n, p_n)_{n \geq 1}$ defined on $(\Omega, \mathcal{A})$ with values in $\mathbb{R}_M^K \times \mathbb{R}_M \times \Delta(K)$ satisfying: (i) $g_1 = g$ a.s., (ii) $(g_n)_{n \geq 1}$ is a martingale adapted to $(\mathcal{F}_n)_{n \geq 1}$, (iii) $\forall n \geq 1, a_{n+1} = a_n$ a.s. or $p_{n+1} = p_n$ a.s., and (iv) $(g_n)_n$ converges a.s. to a random variable g_∞ with values in G .

Let us forget for a while the component of player 2's payoff. A process $(g_n)_n$ satisfying (ii) and (iii) may be called a bi-martingale; it is a martingale such that at every stage, one of the two components remains a.s. constant. So the set G^* can be seen as the set of starting points of converging bi-martingales with limit points in G .

Theorem 11 Hart (1985) Let (a, β) be in $\mathbb{R}^K \times \mathbb{R}$.

$$(a, \beta) \text{ is an equilibrium payoff of } \Gamma(p_0) \\ \Leftrightarrow (a, \beta, p_0) \in G^*.$$

Theorem 11 is too elaborate to be proved here, but let us give a few ideas about the proof. First consider the implication $\Rightarrow$ and fix an equilibrium $\sigma^* = (\sigma^1, \sigma^2)$ of $\Gamma(p_0)$ with payoff (a, β) . The sequence of *a posteriori* $(p_t(\sigma^*))_{t \geq 0}$ is a $\mathbb{P}_{p_0, \sigma^*}$ -martingale. Modify now slightly the time structure so that at each stage, player 1 plays first, and then player 2 plays without knowing the action chosen by player 1. At each half-stage where player 2 plays, his *a posteriori* remains constant. At each half-stage where player 1 plays, the

“expectation of player 1's future payoff” (which can be properly defined) remains constant. Hence, the heuristic apparition of the bimartingale. And since bounded martingale converge, for large stages everything will be fixed and the players will approximately play a nonrevealing equilibrium at a “limit *a posteriori*,” so the convergence will be towards elements of G .

Consider now the converse implication $\Leftarrow$. Let (a, β) be such that $(a, \beta, p_0) \in G^*$ and assume for simplification that the associated bi-martingale (a_n, β_n, p_n) converges in a fixed number N of stages: $\forall n \geq N, (a_n, \beta_n, p_n) = (a_N, \beta_N, p_N) \in G$. One can construct an equilibrium (σ^1, σ^2) of $\Gamma(p_0)$ with payoff (a, β) along the following lines. For each index n , (a_n, β_n) will be an equilibrium payoff of the repeated game with initial probability p_n . Eventually, player 1 will play independently of the state, the *a posteriori* of player 2 will be p_N , and the players will end up playing a nonrevealing equilibrium of the repeated game $\Gamma(p_N)$ with payoff (a_N, β_N) . What should be played before? Since we are in an undiscounted setup, any finite number of stages can be used for communication without influencing payoffs. Let $n < N$ be such that $a_{n+1} = a_n$. To move from (a_n, β_n, p_n) to $(a_n, \beta_{n+1}, p_{n+1})$, player 1 can simply use the splitting lemma (Lemma 1) in order to signal part of the state to player 2. Let now $n < N$ be such that $p_{n+1} = p_n$, so that we want to move from (a_n, β_n, p_n) to $(a_{n+1}, \beta_{n+1}, p_n)$. Player 1 will play independently of the state, and both players will act so as to convexify their future payoffs. This convexification is done through procedures called “jointly controlled lotteries” and introduced in the sixties by Aumann and Maschler (1995), with the following simple and brilliant idea. Imagine that the players have to decide with even probability whether to play the equilibrium $E1$ with payoff (a^1, β^1) or to play the equilibrium $E2$ with payoff (a^2, β^2) . The players may not be indifferent between $E1$ and $E2$, e.g., player 1 may prefer $E1$, whereas player 2 prefers $E2$. They will proceed as follows, with i and i' , respectively, j and j' , denoting two distinct actions of player 1, resp. player 2. Simultaneously and independently, player 1 will select i or i' with probability 1/2, whereas player 2 will behave similarly

with j and j' . $i \begin{pmatrix} j & j' \\ \times & \times \end{pmatrix}_{i'}$. Then the equilibrium

$E1$ will be played if the diagonal has been reached, i.e., if (i, j) or (i', j') has been played, and otherwise the equilibrium $E2$ will be played. This procedure is robust to unilateral deviations: none of the players can deviate and prevent $E1$ and $E2$ to be chosen with probability $1/2$. In general, jointly controlled lotteries are procedures allowing to select an alternative among a finite set according to a given probability (think of binary expansions if necessary), in a way which is robust to deviations by a single player. S. Hart has precisely shown how to combine steps of signaling and jointly controlled lotteries to construct an equilibrium of $\Gamma_\infty(p_0)$ with payoff (a, β) .

Biconvexity and Bimartingales

The previous analysis has led to the introduction and study of biconvexity phenomena. The reference here is Aumann and Hart (1986). Let X and Y be compact convex subsets of Euclidean spaces, and let $(\Omega, \mathcal{F}, \mathcal{P})$ be an atomless probability space.

Definition 13 A subset B of $X \times Y$ is biconvex if for every x in X and y in Y , the sections $B_x = \{y' \in Y, (x, y') \in B\}$ and $B_y = \{x' \in X, (x', y) \in B\}$ are convex. If B is biconvex, a mapping $f: B \rightarrow \mathbb{R}$ is called biconvex if for each $(x, y) \in X \times Y$, $f(\cdot, y)$ and $f(x, \cdot)$ are convex.

As in the usual convexity case, we have that if f is biconvex, then for each α in $\mathbb{R}$, the set $\{(x, y) \in B, f(x, y) \leq \alpha\}$ is biconvex.

Definition 14 A sequence of random variables $Z_n = (X_n, Y_n)_{n \geq 1}$ with values in $X \times Y$ is called a bimartingale if:

- (1) There exists an increasing sequence $(\mathcal{F}_n)_{n \geq 1}$ of finite sub- σ -algebra of $\mathcal{F}$ such that $(Z_n)_n$ is a $(\mathcal{F}_n)_{n \geq 1}$ -martingale.
- (2) $\forall n \geq 1, X_n = X_{n+1}$ a.s. or $Y_n = Y_{n+1}$ a.s.
- (3) Z_1 is a.s. constant.

Notice that $(Z_n)_{n \geq 1}$ being a bounded martingale, it converges almost surely to a limit Z_∞ .

Definition 15 Let A be a measurable subset of $X \times Y$.


$A^* = \{z \in X \times Y, \text{ there exists a bimartingale } (Z_n)_{n \geq 1} \text{ converging to a limit } Z_\infty \text{ such that } Z_\infty \in A \text{ a.s. and } Z_1 = z \text{ a.s.}\}$.

One can show that any atomless probability space $(\Omega, \mathcal{F}, \mathcal{P})$, or any product of convex compact spaces $X \times Y$ containing A , induces the same set A^* . One can also substitute condition (2) by: $\forall n \geq 1, (X_n = X_{n+1} \text{ or } Y_n = Y_{n+1})$ a.s. Notice that without condition (2), the set A^* would just be the convex hull of A .

We always have $A \subset A^* \subset \text{conv}(A)$, and these inclusions can be strict. For example, if $X = Y = [0, 1]$ and $A = \{(0, 0), (1, 0), (0, 1)\}$, it is possible to show that $A^* = \{(x, y) \in [0, 1] \times [0, 1], x = 0 \text{ or } y = 0\}$. A^* always is biconvex and thus contains $\text{biconv}(A)$, which is defined as the smallest biconvex set which contains A . The inclusion $\text{biconv}(A) \subset A^*$ can also be strict, as shown by the following example:

Example 5 Put $X = Y = [0, 1]$, $v_1 = (1/3, 0)$, $v_2 = (0, 2/3)$, $v_3 = (2/3, 1)$, $v_4 = (1, 1/3)$, $w_1 = (1/3, 1/3)$, $w_2 = (1/3, 2/3)$, $w_3 = (2/3, 2/3)$ et $w_4 = (2/3, 1/3)$, and $A = \{v_1, v_2, v_3, v_4\}$ (Fig. 5).

A is biconvex, so $A = \text{biconv}(A)$. Consider now the following Markov process $(Z_n)_{n \geq 1}$, with $Z_1 = w_1$. If $Z_n \in A$, then $Z_{n+1} = Z_n$. If $Z_n = w_i$ for some i , then $Z_{n+1} = w_{i+1 \pmod{4}}$ with probability $1/2$, and $Z_{n+1} = v_i$ with probability $1/2$. $(Z_n)_n$ is a



Repeated Games with Incomplete Information, Fig. 5 The “four frogs” example of Aumann and Hart: $A^* \neq \text{biconv}(A)$

bimartingale converging a.s. to a point in A , hence $w_1 \in A^* \setminus \text{biconv}(A)$.

We now present a geometric characterization of the set A^* and assume here that A is closed. For each biconvex subset B of $X \times Y$ containing A , we denote by $\text{nsc}(B)$ the set of elements of B which cannot be separated from A by a continuous bounded biconvex function on A . More precisely, $\text{nsc}(B) = \{z \in B, \forall f: B \rightarrow \mathbb{R} \text{ bounded biconvex, and continuous on } A, f(z) \leq \sup\{f(z'), z' \in A\}\}$.

Theorem 12 Aumann and Hart (1986): A^* is the largest biconvex set B containing A such that $\text{nsc}(B) = B$.

Let us now come back to repeated games and to the notations of subsection “[Characterization of Equilibrium Payoffs](#).” To be precise, we need to add the component of player 2’s payoff and consequently to slightly modify the definitions. G is closed in $\mathbb{R}_M^K \times \mathbb{R}_M \times \Delta(K)$. For $B \subset \mathbb{R}_M^K \times \mathbb{R}_M \times \Delta(K)$, B is biconvex if for each a in $\mathbb{R}_M^K$ and for each p in $\Delta(K)$, the sections $\{(\beta, p'), (a, \beta, p') \in B\}$ and $\{(a', \beta), (a', \beta, p) \in B\}$ are convex. A real function f defined on a biconvex set B is said to be biconvex if $\forall a, \forall p, f(a, \cdot, p)$ and $f(\cdot, \cdot, p)$ are convex.

Theorem 13 Aumann and Hart (1986): G^* is the largest biconvex set B containing G such that: $\forall z \in B, \forall f: B \rightarrow \mathbb{R}$ bounded biconvex, and continuous on $A, f(z) \leq \sup\{f(z'), z' \in G\}$.

Nonobservable Actions

We now consider the case where, as in the general definition of section “[Definition of the Subject and Its Importance](#),” there is a signaling function $q: K \times A \rightarrow \Delta(U)$ giving the distributions of the signals received by the players as a function of the state of nature and the action profile just played. The particular case where $q(k, a)$ does not depend on k is called *state independent signaling*. The previous models correspond to the particular case of perfect observation, where the signals received by the players exactly reveal the action profile played.

Theorem 1 has been generalized (Aumann and Maschler 1995) to the general case of signaling function. We keep the notations of section “[The Standard Model of Aumann and Maschler](#).” Given a mixed action $x \in \Delta(I)$, an action j in J and a state k , we denote by $Q(k, x, j)$ the marginal distribution on U^2 of the law $\sum_{i \in I} x(i) q(k, i, j)$, i.e., $Q(k, x, j)$ is the law of the signal received by player 2 if the state is k , player 1 uses x and player 2 plays j . The set of nonrevealing strategies of player 1 is then defined as: $NR(p) = \left\{x = (x^k)_{k \in K} \in \Delta(I)^K, \forall k \in K, \forall k' \in K \text{ s.t. } p^k p^{k'} > 0, \forall j \in J, Q(k, x^k, j) = Q(k', x^{k'}, j)\right\}$. If the initial probability is p and player 1 plays a strategy x in $NR(p)$ (i.e., plays x^k if the state is k), the *a posteriori* of player 2 will remain a.s. constant: player 2 can deduce no information on the selected state k . The value of the non-revealing game becomes:

$$\begin{aligned} u(p) &= \max_{x \in NR(p)} \min_{y \in \Delta(J)} \sum_{k \in K} p^k G^k(x^k, y) \\ &= \min_{y \in \Delta(J)} \max_{x \in NR(p)} \sum_{k \in K} p^k G^k(x^k, y), \end{aligned}$$

where $G^k(x^k, y) = \sum_{i, j} x^k(i) y(j) G^k(i, j)$, and the convention $u(p) = -\infty$ if $NR(p) = \emptyset$. Theorem 1 perfectly extends here: The repeated game with initial probability p has a uniform value given by $\text{cav}u(p)$.

The explicit construction of an optimal strategy of player 2 (see section “[Back to the Standard Model](#)” here) has also been generalized to the general signaling case (see Kohlberg 1975; Mertens et al. 1994, part B, p.234 for random signals).

Regarding zero-sum games with lack of information on both sides, the results of section “[Zero-Sum Games with Lack of Information on Both Sides](#)” have been generalized to the case of state independent signaling (see Mertens 1972; Mertens and Zamir 1971, 1977). Attention has been paid to the speed of convergence of the value function $(v_T)_T$, and bounds are identical for both models of lack of information on one side and on both sides, if we assume state independent signaling: this speed is of order $1/T^{1/2}$ for games with perfect observation and of order $1/T^{1/3}$ for games with signals (these orders are optimal, both for lack of information on one side

and lack of information on both sides, see (Zamir 1971, 1973). For state-dependent signaling and lack of information on one side, it was shown by Mertens (1998) that the convergence occurs with worst case error $\sim (\ln n/n)^{1/3}$.

A particular class of zero-sum repeated games with state dependent signaling has been studied (games with no signals, see (Mertens and Zamir 1976b; Sorin 1989; Waternaux 1983). In these games, the state k is first selected according to a known probability and is not announced to the players; then after each stage both players receive the same signal which is either “nothing” or “the state is k .” It was shown that the maxmin and the minmax may differ, although $\lim_T v_T$ always exists.

In nonzero-sum repeated games with lack of information on one side, the existence of “joint plan” equilibria have been generalized to the case of state independent signaling (Renault 2000) and more generally to the case where “player 1 can send non revealing signals to player 2” (Simon et al. 2002). The existence of a uniform equilibrium in the general signaling case is still an open question (see Simon et al. 2008).

Advances

1. Zero-sum games with lack of information on one and a half side

In games with lack of information on one side, it is important that player 1 knows not only the selected state k , but also the *a priori* p . Sorin and Zamir (1985) provide an example of a game with lack of information on “one and a half” side with no uniform value. More precisely, in this example nature first chooses p in $\{p_1, p_2\}$ according to a known probability and announces p to player 2 only; then k is selected according to p , and announced to player 1 only; finally the matrix game G^k is played.

2. v_T and v_λ as a function of p

For games with lack of information on one side, the value function v_T is a concave piecewise linear function of the initial probability p (see Ponsard and Sorin 1980 for more generality). On the

contrary, the discounted value v_λ can be quite a complex function of p : in Example 2 of section “Definition of the Subject and Its Importance,” Mayberry (1967) has proved that for $2/3 < \lambda < 1$, v_λ is, at each rational value of p , nondifferentiable.

3. $\lim_T \sqrt{T}(v_T(p) - \text{cavu}(p))$ and the normal distribution

Convergence of the value functions $(v_T)_T$ and $(v_\lambda)_\lambda$ has been widely studied. We have already mentioned the speed of convergence in section “Non-Observable Actions,” but much more can be said.

Example 6 Standard model of lack of information

on one side and observable actions. $K = \{a, b\}$, $G^a = \begin{pmatrix} 3 & -1 \\ -3 & 1 \end{pmatrix}$ and $G^b = \begin{pmatrix} 2 & -2 \\ -2 & 2 \end{pmatrix}$. One can show (Mertens and Zamir 1976a) that for each $p \in [0, 1]$, viewed as the initial probability of state a , the sequence $\sqrt{T} v_T(p)$ converges to $\varphi(p)$, where $\varphi(p) = \frac{1}{\sqrt{2\pi}} e^{-x_p^2/2}$, and x_p satisfies $\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x_p} e^{-x^2/2} dx = p$. So the limit of $\sqrt{T} v_T(p)$ is the standard normal density function evaluated at its p -quantile.

The apparition of the normal distribution is by no way an isolated phenomenon, but rather an important property of some repeated games (de Meyer 1996a, b, 1998, 1999; de Meyer and Moussa Saley 2003, ...).

4. The dual game

B. de Meyer introduced the notion of “dual game” (see the previous references and also de Meyer and Marino 2005; de Meyer and Rosenberg 1999; Laraki 2002; Rosenberg 1998). Let us now illustrate this on the standard model of section “The Standard Model of Aumann and Maschler.”

Let z be a parameter in $\mathbb{R}^K$. In the dual game $\Gamma_T^*(z)$, player 1 first secretly chooses the state k . Then at each stage $t \leq T$, the players choose as usual actions it and jt which are announced before proceeding to the next stage. With time horizon T ,

player 1's payoff finally is $\frac{1}{T} \sum_{t=1}^T G^k(i_t, j_t) - z^k$. This player is thus now able to fix the state equal to k , but has to pay z^k for it. It can be shown that the T -stage dual game $\Gamma_T^*(z)$ has a value $w_T(z)$. w_T is convex and is linked to the value of the primal game by the conjugate formula:

$$w_T(z) = \max_{p \in \Delta(K)} (v_T(p) - \langle p, z \rangle), \text{ and}$$

$$v_T(p) = \inf_{z \in \mathbb{R}^K} (w_T(z) + \langle p, z \rangle).$$

And $(w_T)T$ satisfies the dual recursive formula:

$$w_{T+1}(z) = \min_{y \in \Delta(J)} \max_{i \in I} \frac{T}{T+1} w_T \left(\frac{T+1}{T} z - \frac{1}{T} \sum_{j \in J} y_j (G^k(i, j))_k \right)$$

There are also strong relations between the optimal strategies of the players in the primal and dual games, and this gives a way to compute recursively optimal strategies of the uninformed player in the finite game (see also Heuer 1992 on this topic).

5. Approachability

Blackwell's approachability theorem has been extended to infinite dimensional spaces by Lehrer (2003a). As we saw in Theorem 6, approachability theory has strong links with the existence of no-regret strategies (first studied in Hart and Mas-Colell 2000), see also Cesa-Bianchi et al. (2006); Foster and Vohra (1999); Hart (2005); Lehrer (2003b); Rustichini (1999) and the book Cesa-Bianchi and Lugosi (2006), but also with convergence of simple procedures to the set of correlated equilibria (Hart and Mas-Colell 2000) and calibration (Foster 1999; Lehrer 2001). The links between merging, reputation phenomena, and repeated games with incomplete information have been studied in (Sorin 1997), where several existing results are unified. And no-regret and approachability have also been studied when the players have bounded computational capacities (finite automata, bounded recall strategies) (Lehrer and Solan 2003, 2006).

6. Markov chain games with lack of information

In Renault (2006), the standard model of lack of information, as well as the proof of Theorem 1, is generalized to the case where the state is not fixed at the beginning of the game but evolves according to a Markov chain uniquely observed by player 1 (see also Neyman (2008) for nonobservable actions). The limit value is however difficult to compute, as shown by the following example from Renault (2006): $K = \{a, b\}$, the payoff matrices are $G^a = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$ and $G^b = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$, the initial probability is $(1/2, 1/2)$, and the state evolves according to the Markov chain $M = \begin{pmatrix} \alpha & 1-\alpha \\ 1-\alpha & \alpha \end{pmatrix}$ with parameter α . If $\alpha = 1$ this is Example 2, and the limit value is $1/4$ by Theorem 1.

For $\alpha \in [1/2, 2/3]$, the limit value is $\frac{\alpha}{4\alpha-1}$ (Marino 2005) for $\alpha = 2/3$, (Hörner et al. 2010).

For $\alpha \in [2/3, .719]$, the limit value v satisfies (Bressaud and Quas 2006): $\frac{1}{v} = u_0 + u_0 u_1 + u_0 u_1 u_2 + \dots$, where (u_n) is defined by $u_0 = 1$ and $u_{n+1} = \max\{\psi(u_n), 1 - \psi(u_n)\}$, with $\psi(u) = 3\alpha - 1 - \frac{2\alpha-1}{u}$. What is the value for $\alpha = 0.9$?

In Markov chain games with lack of information on both sides, each player privately observes his state variable, and both state variables follow exogenous and independent Markov chains. For such games, the existence of the limit value $\lim_T v_T = \lim_{\lambda} v_{\lambda}$ has been proved in Gensbittel and Renault (2015). In the case of recurrent and aperiodic chains, the limit value is identified as the unique solution of the Mertens-Zamir system of Theorem 8, with an appropriate nonrevealing function $\hat{u}$ instead of u , corresponding to the limit value of the auxiliary dynamic game where each player is restricted to play strategies that reveal no information on the recurrence class of his own state (such a nonrevealing function was already considered in Renault (2006)).

In the nonzero-sum context, dynamic sender-receiver games are Markov chain games with lack of information on one side, where the payoffs only depend on the actions of the uninformed player. The set of equilibrium payoffs has been characterized under a homothety (random shocks) assumption on the Markov chain (Renault et al. 2013).

7. Extension to zero-sum dynamic games with state process controlled and observed by player 1

It is known since (Sorin 1984a) that the uniform value may not exist in general for stochastic games with lack of information on one side on the payoff matrices (where the payoff matrices of the stochastic game to be played are first randomly selected and announced to player 1 only). Rosenberg et al. (2004) studied stochastic games with a single controller and lack of information on one side on the payoff matrices, showing the existence of the uniform value if the informed player controls the transition, and providing a counter-example if the uninformed player controls the transitions. One can also consider the model of general repeated games with an informed controller (Renault 2012), generalizing the model of Markov chain games with lack of information on one side), i.e., dynamic games with finitely many states, actions and signals, and state processes controlled and observed by player 1.

A general repeated game is given by: 5 non empty finite sets: a set of states or parameters K , a set I of actions for player 1, a set J of actions for player 2, a set C of signals for player 1, and a set D of signals for player 2, an initial distribution $\pi \in \Delta(K \times C \times D)$, a payoff function $g: K \times I \times J \rightarrow [0, 1]$ for player 1, and a transition function $q: K \times I \times J \rightarrow \Delta(K \times C \times D)$. The progress of the game is the following: Initially, (k_1, c_1, d_1) is selected according to π , player 1 learns c_1 and player 2 learns d_1 . Then simultaneously player 1 chooses i_1 in I and player 2 chooses j_1 in J , and the payoff for player 1 at stage 1 is $g(k_1, i_1, j_1)$, etc. At any stage $t \geq 2$, (k_t, c_t, d_t) is selected according to $q(k_{t-1}, i_{t-1}, j_{t-1})$, player 1 learns c_t and player 2 learns d_t . Simultaneously, player 1 chooses i_t in I and player 2 chooses j_t in J . The stage payoffs are $g(k_t, i_t, j_t)$ for player 1 and the opposite for player 2, and the play proceeds to stage $t + 1$.

In repeated games with an informed controller, it is moreover assumed that:

- 1) Player 1 is fully informed, in the sense that π and q are such that he can always deduce the state and player 2's signal from his own signal.
- 2) Player 1 controls the transition, in the sense that the marginal $\bar{q}$ of the transition q on $K \times D$ does not depend on player 2's action.

In this setup, one can prove (Renault 2012) the existence of the uniform value $v^*(\pi)$, satisfying:

$$v^*(\pi) = \inf_{n \geq 1} \sup_{m \geq 0} v_{m,n}(\pi) = \sup_{m \geq 0} \inf_{n \geq 1} v_{m,n}(\pi).$$

where $v_{m,n}(\pi)$ is the value of the game with payoff $\mathbb{E}_{\pi, \sigma, \tau} \frac{1}{n} \left(\sum_{t=m+1}^{m+n} g_t \right)$, g_t being the payoff of stage t .

Moreover, one can prove for such games the existence of the stronger notion of “general uniform value.” Let us first define the values $v_\theta(\pi)$ of the dynamic game with payoff $\gamma_\theta(\pi, \sigma, \tau) = \mathbb{E}_{\pi, \sigma, \tau} \left(\sum_{t \geq 1} \theta_t g_t \right)$, where θ is an evaluation $(\theta_t)_{t \geq 1}$ with nonnegative weights satisfying $\sum_{t \geq 1} \theta_t = 1$, and total variation denoted by $TV(\theta) = \sum_t |\theta_{t+1} - \theta_t|$. And $v^*(\pi)$ is the general uniform value of the game with initial probability π if for each $\varepsilon > 0$ one can find $\alpha > 0$ and a couple of strategies σ^* and τ^* such that for all evaluations θ with $TV(\theta) \leq \alpha$:

$$\forall \tau, \gamma_\theta(\pi, \sigma^*, \tau) \geq v^*(\pi) - \varepsilon \quad \text{and} \quad \forall \sigma, \gamma_\theta(\pi, \sigma, \tau^*) \leq v^*(\pi) + \varepsilon.$$

Considering only Cesaro-evaluations (i.e., of the type $\theta_t = 1/n$ for $t \leq n$, $= 0$ for $t > n$ for some n) recovers our Definition 1. Renault and Venel (2017) introduce a new distance (compatible with the weak topology) on the belief space $\Delta(\Delta(K))$ of Borel probabilities over the simplex $X = \Delta(K)$ and prove the existence of the general uniform value in general repeated games with an informed controller. Clearly, the values only depend on player 2's belief p on the initial state, and the limit value v^* can be characterized as:

$$\forall p \in X, v^*(p) = \inf \{ w(p), w: \Delta(X) \rightarrow [0, 1] \text{ affine } C^0 \text{ s.t.} \}$$

1. $\forall p' \in X, w(p') \geq \sup_{a \in \Delta(I)^K} w(\bar{q}(p', a))$
2. $\forall (z, y) \in RR, w(z) \geq y$.

where $\bar{q}(p, a) = \sum_{k \in K} p^k \bar{q}(k, a^k) \in \Delta(K \times D)$ gives the marginal of q on $K \times D$, and $RR = \{(z, y) \in \Delta(X) \times [0, 1], \text{ there exists } a: X \rightarrow \Delta(I)^K \text{ measurable s.t. } \int_{p \in X} \bar{q}(p, a(p)) dz(p) = z \text{ and } \int_{p \in X} \min_{j \in J} (\sum_{k \in K} p^k g(k, a^k, j)) dz(p) = y\}$ can be seen as the set of invariant measures and

associated payoffs. In the standard model of Aumann and Maschler, (1) is equivalent to w being a concave function on $\Delta(K)$ and (2) is equivalent to w being not lower than the non-revealing function u : so v^* is the smallest concave function above u , and we recover the *cavu* theorem (Theorem 1).

Finally, the existence of the uniform value has been generalized to the case where Player 1 controls the transitions and is more informed than player 2 (but player 1 does not necessarily observe the current state) in Gensbittel et al. (2014).

8. Symmetric information

Another model deals with the symmetric case, where the players have an incomplete, but identical, knowledge of the selected state. After each stage, they receive the same signal, which may depend on the state. A. Neyman and S. Sorin have proved the existence of equilibrium payoffs in the case of two players (see Neyman and Sorin 1998, the zero-sum case being solved in Forges 1982; Kohlberg and Zamir 1974).

This result does not extend to the case where the stage evolves from stage to stage, i.e., to stochastic games with incomplete information. In the zero-sum symmetric information case where at the end of each stage, the players observe both actions but receive no further information on the current state (*hidden stochastic games*), B. Ziliotto provided in his PhD thesis an example where $\lim_T v_T$ and $\lim_\lambda v_\lambda$ may fail to exist (Ziliotto 2016).

One can also consider zero-sum general repeated games with payoffs defined by the expectation of a Borel function over plays. In the public case where the players have the same information at the end of every stage, the value exists (Gimbert et al. 2016).

9. Continuous-time approach

A continuous time approach can also be used to prove convergence results in general zero-sum repeated games, and in particular Theorem 7, embedding the discrete repeated game into a continuous time game and using viscosity solution

tools (Cardaliaguet et al. 2012). A generalization of the *cavu* theorem (Theorem 1) to infinite action spaces and partial information can be found in Gensbittel (2015), using a probabilistic method based on martingales and a functional method based on approximation schemes for viscosity solutions of Hamilton Jacobi equations.

10. The operator approach for zero-sum games

Repeated games with incomplete information, as well as stochastic games, can also be studied in a functional analysis setup called the operator approach. This general approach is based on the study of the recursive formula (Laraki 2001b; Rosenberg and Sorin 2001; Sorin 2002).

11. Uncertain duration

One can consider zero-sum repeated games with incomplete information on both sides and *uncertain duration*. In these games, the payoff to the players is the sum of their stage payoffs, up to some stopping time θ which may depend on plays, divided by the expectation of θ . Theorem 8 here generalizes to the case of *public* uncertain duration process (as $\mathbb{E}(\theta) \rightarrow \infty$, with a convergence in $O\left(1/\mathbb{E}(\sqrt{\theta})\right)$), see Neyman and Sorin (2010).

The situation is different if one allows for private uncertain duration processes: any number between the $\max_{\min} \text{cav}_I \text{vex}_I \mu(p, q)$ and the $\min_{\max} \text{vex}_I \text{cav}_I \mu(p, q)$ is the value of a long finitely repeated game Γ_T where players' information about the uncertain number of repetitions T is asymmetric (Neyman 2012).

12. Frequent actions

One can consider a repeated game with incomplete information and fixed discount factor, where the time span between two consecutive stages is $1/n$. In the context of zero-sum Markov chain games of lack of information on one side, Cardaliaguet et al. (2016) show the existence of a limit value when n goes to $+\infty$; this value is characterized through an auxiliary stochastic optimization problem and, independently, as the solution of an Hamilton-Jacobi equation.

13. Repeated market games with incomplete information

De Meyer and Moussa Saley studied the modelization via Brownian motions in financial models (de Meyer and Moussa Saley 2003). They introduced a marked game based on a repeated game with lack of information on one side and showed the endogenous apparition of a Brownian motion (see also de Meyer and Marino 2004 for incomplete information on both sides, and de Meyer 2010).

14. Cheap-talk and communication

In the nonzero-sum setup of section “[Non Zero-Sum Games with Lack of Information on One Side](#),” it is interesting to study the number of communication stages which is needed to construct the different equilibria. This number is linked with the convergence of the associated bimartingales (see Aumann and Hart 1986; Forges 1984, 1990; Aumann and Maschler 1995). Let us mention also that F. Forges (1988) gave a similar characterization of equilibrium payoffs, for a larger notion of equilibria called communication equilibria (see also Forges 1985 for correlated equilibria). Amitai (1996b) studied the set of equilibrium payoffs in case of lack of information on both sides. Aumann and Hart (2003) characterized the equilibrium payoffs in two player games with lack of information on one side when long, payoff-irrelevant, preplay communication is allowed (see Amitai 1996a for incomplete information on both sides).

15. Known own payoffs

The particular nonzero-sum case where each player knows his own payoffs is particularly worthwhile studying. In the two-player case with lack of information on one side, this amounts to say that player 2's payoffs do not depend on the selected state. In this case, Shalev (1994) showed that any equilibrium payoff can be obtained as the payoff of an equilibrium which is completely revealing. This result generalizes to the nonzero-sum case of lack of information of both sides (see

the unpublished manuscript Koren 1992), but uniform equilibria may fail to exist even though both players know their own payoffs.

16. More than 2 players

Few papers study the case of more than 2 players. The existence of uniform equilibrium has been studied for 3 players and lack of information on one side (Renault 2001a), and in the case of two states of nature it appears that a completely revealing equilibria, or a joint plan equilibria by one of the informed players, always exists. Concerning n -player repeated games with incomplete information and signals, several papers study how the initial information can be strategically transmitted, independently of the payoffs (Renault 2001b; Renault et al. 2014; Renault and Tomala 2004, 2008), with cryptographic considerations. As an application, the existence of completely revealing equilibria, i.e., equilibria where each player eventually learns the state with probability one, is obtained in particular cases (see also Hörner et al. 2011 for the related notion of “belief-free” equilibria).

17. Perturbations of repeated games with complete information

Repeated games with incomplete information have been used to study perturbations of repeated games with complete information (see Cripps and Thomas 2003; Fudenberg and Maskin 1986) for Folk theorem-like results (Aumann and Sorin 1989), for enforcing cooperation in games with a Pareto dominant outcome, and (Israeli 2010) for a perturbation with known own payoffs). The case where the players have different discount factors has also been investigated (Cripps and Thomas 2003; Lehrer and Yariv 1999).

Future Directions

Several open problems are well formulated and deserve attention. Does a uniform equilibrium always exist in two-player repeated games with lack of information on one side and general

signaling or in n -player repeated games with lack of information on one side? Does the limit value always exist in zero-sum repeated games with incomplete information and signals? More conceptually, one should look for classes of n -player repeated games with incomplete information which allow for the existence of equilibria, and/or for a tractable description of equilibrium payoffs (or at least of some of these payoffs). Regarding applications, there is certainly a lot of room in the vast fields of financial markets, cryptography, learning, and sequential decision problems.

Bibliography

Primary Literature

- Amitai M (1996a) Cheap-talk with incomplete information on both sides. PhD thesis, The Hebrew University of Jerusalem. <http://ratio.huji.ac.il/dp/dp90.pdf>
- Amitai M (1996b) Repeated games with incomplete information on both sides. PhD thesis, The Hebrew University of Jerusalem. <http://ratio.huji.ac.il/dp/dp105.pdf>
- Aumann RJ (1964) Mixed and behaviour strategies in infinite extensive games. In: Dresher M, Shapley LS, Tucker AW (eds) *Advances in game theory*. Annals of Mathematics Study 52. Princeton University Press, pp 627–650
- Aumann RJ, Hart S (1986) Bi-convexity and bi-martingales. *Israel J Math* 54:159–180
- Aumann RJ, Hart S (2003) Long cheap talk. *Econometrica* 71:1619–1660
- Aumann RJ, Sorin S (1989) Cooperation and bounded recall. *Games Econom Behav* 1:5–39
- Blackwell D (1956) An analog of the minmax theorem for vector payoffs. *Pac J Math* 65:1–8
- Bressaud X, Quas A (2006) Dynamical analysis of a repeated game with incomplete information. *Math Oper Res* 31:562–580
- Cardaliaguet P, Rainer C, Rosenberg D, Vieille N (2016) Markov games with frequent actions and incomplete information? The limit case. *Math Oper Res* 41:49–71
- Cardaliaguet P, Laraki R, Sorin S (2012) A continuous time approach for the asymptotic value in two-person zero-sum repeated games. *SIAM J Control Optim* 50:1573–1596
- Cesa-Bianchi N, Lugosi G (2006) *Prediction, learning and games*. Cambridge University Press, Cambridge
- Cesa-Bianchi N, Lugosi G, Stoltz G (2006) Regret minimization under partial monitoring. *Math Oper Res* 31:562–580
- Cripps MW, Thomas JP (2003) Some asymptotic results in discounted repeated games of one-sided incomplete information. *Math Oper Res* 28:433–462
- de Meyer B (1996a) Repeated games and partial differential equations. *Math Oper Res* 21:209–236
- de Meyer B (1996b) Repeated games, duality and the central limit theorem. *Math Oper Res* 21:237–251
- de Meyer B (1998) The maximal variation of a bounded martingale and the central limit theorem. *Ann Inst Henri Poincaré Probab Stat* 34:49–59
- de Meyer B (1999) From repeated games to Brownian games. *Annales de l'Institut Henri Poincaré, Probabilités et statistiques* 35:1–48
- de Meyer B (2010) Price dynamics on a stock market with asymmetric information. *Games Econom Behav* 69:42–71
- de Meyer B, Marino A (2004) Repeated market games with lack of information on both sides. DP 2004.66, MSE Université Paris I
- de Meyer B, Marino A (2005) Duality and optimal strategies in the finitely repeated zero-sum games with incomplete information on both sides. DP 2005.27, MSE Université Paris I
- de Meyer B, Moussa Saley H (2003) On the strategic origin of Brownian motion in finance. *Int J Game Theory* 31:285–319
- de Meyer B, Rosenberg D (1999) “Cavu” and the dual game. *Math Oper Res* 24:619–626
- Forges F (1982) Infinitely repeated games of incomplete information: symmetric case with random signals. *Int J Game Theory* 11:203–213
- Forges F (1984) A note on Nash equilibria in repeated games with incomplete information. *Int J Game Theory* 13:179–187
- Forges F (1985) Correlated equilibria in a class of repeated games with incomplete information. *Int J Game Theory* 14:129–149
- Forges F (1988) Communication equilibria in repeated games with incomplete information. *Math Oper Res* 13:191–231
- Forges F (1990) Equilibria with communication in a job market example. *Q J Econ* 105:375–398
- Foster D (1999) A proof of calibration via Blackwell’s approachability theorem. *Games Econom Behav* 29:73–78
- Foster D, Vohra R (1999) Regret in the on-line decision problem. *Games Econom Behav* 29:7–35
- Fudenberg D, Maskin E (1986) The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* 54:533–554
- Gensbittel F (2015) Extensions of the Cav(u) theorem for repeated games with one-sided information. *Math Oper Res* 40(1):80–104
- Gensbittel F, Renault J (2015) The value of Markov Chain Games with incomplete information on both sides. *Math Oper Res* 40(4):820–841
- Gensbittel F, Oliu-Barton M, Venel X (2014) Existence of the uniform value in repeated games with a more informed controller. *J Dynam Games* 1(3):411–445
- Gimbert H, Renault J, Sorin S, Zielonka W (2016) On values of repeated games with signals. *Ann Appl Probab* 26:402–424
- Harsanyi J (1967–68) Games with incomplete information played by ‘Bayesian’ players, parts I–III. *Manag Sci* 8:159–182, 320–334, 486–502
- Hart S (1985) Nonzero-sum two-person repeated games with incomplete information. *Math Oper Res* 10:117–153

- Hart S (2005) Adaptive Heuristics. *Econometrica* 73:1401–1430
- Hart S, Mas-Colell A (2000) A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68:1127–1150
- Heuer M (1992) Optimal strategies for the uninformed player. *Int J Game Theory* 20:33–51
- Horner J, Lovo S, Tomala T (2011) Belief-free equilibria in games with incomplete information: characterization and existence. *J Econ Theory* (5):1770–1795
- Horner J, Rosenberg D, Solan E, Vieille N (2010) On a Markov game with one-sided incomplete information. *Oper Res* 58:1107–1115
- Israeli E (2010) Sowing doubt optimally in two-person repeated games. *Games Econ Behav* 28:203–216. 1999
- Kohlberg E (1975) Optimal strategies in repeated games with incomplete information. *Int J Game Theory* 4:7–24
- Kohlberg E, Zamir S (1974) Repeated games of incomplete information: the symmetric case. *Ann Stat* 2:40–41
- Koren G (1992) Two-person repeated games where players know their own payoffs, master thesis, Tel-Aviv University. <http://www.ma.huji.ac.il/hart/papers/koren.pdf>
- Kuhn HW (1953) Extensive games and the problem of information. In: Kuhn and Tucker (eds) *Contributions to the theory of games*, vol II. *Annals of Mathematical Studies* 28. Princeton University Press, pp 193–216
- Laraki R (2001a) Variational inequalities, system of functional equations and incomplete information repeated games. *SIAM J Control Optim* 40:516–524
- Laraki R (2001b) The splitting game and applications. *Int J Game Theory* 30:359–376
- Laraki R (2002) Repeated games with lack of information on one side: the dual differential approach. *Math Oper Res* 27:419–440
- Lehrer E (2001) Any inspection is manipulable. *Econometrica* 69:1333–1347
- Lehrer E (2003a) Approachability in infinite dimensional spaces. *Int J Game Theory* 31:253–268
- Lehrer E (2003b) A wide range no-regret theorem. *Games Econ Behav* 42:101–115
- Lehrer E, Solan E (2003) No regret with bounded computational capacity, *DP 1373*, Center for Mathematical Studies in Economics and Management Science, Northwestern University
- Lehrer E, Solan E (2006) Excludability and bounded computational capacity. *Math Oper Res* 31:637–648
- Lehrer E, Yariv L (1999) Repeated games with lack of information on one side: the case of different discount factors. *Math Oper Res* 24:204–218
- Marino A (2005) The value of a particular Markov chain game. Chapters 5 and 6, PhD thesis, Université Paris I, 2005. <http://alexandre.marino.free.fr/theseMarino.pdf>
- Mayberry J-P (1967) Discounted repeated games with incomplete information, Report of the U.S. Arms control and disarmament agency, ST116, chapter V, *Mathematica*, Princeton, pp 435–461
- Mertens J-F (1972) The value of two-person zero-sum repeated games: the extensive case. *Int J Game Theory* 1:217–227
- Mertens J-F (1998) The speed of convergence in repeated games with incomplete information on one side. *Int J Game Theory* 27:343–357
- Mertens J-F, Zamir S (1971) The value of two-person zero-sum repeated games with lack of information on both sides. *Int J Game Theory* 1:39–64
- Mertens J-F, Zamir S (1976a) The normal distribution and repeated games. *Int J Game Theory* 5:187–197
- Mertens J-F, Zamir S (1976b) On a repeated game without a recursive structure. *Int J Game Theory* 5:173–182
- Mertens J-F, Zamir S (1977) A duality theorem on a pair of simultaneous functional equations. *J Math Anal Appl* 60:550–558
- Mertens J-F, Zamir S (1985) Formulation of Bayesian analysis for games with incomplete information. *Int J Game Theory* 14:1–29
- Neyman A (2008) Existence of optimal strategies in Markov games with incomplete information. *Int J Game Theory* 37:581–596
- Neyman A (2012) The value of two-person zero-sum repeated games with incomplete information and uncertain duration. *Int J Game Theory* 41:95–207
- Neyman A, Sorin S (1998) Equilibria in repeated games with incomplete information: the general symmetric case. *Int J Game Theory* 27:201–210
- Neyman A, Sorin S (2010) Repeated games with public uncertain duration processes. *Int J Game Theory* 39:29–52
- Ponssard JP, Sorin S (1980) The LP formulation of finite zero-sum games with incomplete information. *Int J Game Theory* 9:99–105
- Renault J (2000) 2-player repeated games with lack of information on one side and state independent signaling. *Math Oper Res* 4:552–572
- Renault J (2001a) 3-player repeated games with lack of information on one side. *Int J Game Theory* 30:221–246
- Renault J (2001b) Learning sets in state dependent signaling game forms: a characterization. *Math Oper Res* 26:832–850
- Renault J (2006) The value of Markov chain games with lack of information on one side. *Math Oper Res* 31:490–512
- Renault J (2012) The value of repeated games with an informed controller. *Math Oper Res* 37:154–179
- Renault J, Tomala T (2004) Learning the state of nature in repeated games with incomplete information and signals. *Games Econ Behav* 47:124–156
- Renault J, Tomala T (2008) Probabilistic reliability and privacy of communication using multicast in general neighbor networks. *J Cryptol* 21(2):250–279
- Renault J, Venel X (2017) A distance for probability spaces, and long-term values in Markov decision processes and repeated games. *Math Oper Res* 42(2):349–376
- Renault J, Solan E, Vieille N (2013) Dynamic sender-receiver games. *J Econ Theory* 148:502–534
- Renault J, Renou L, Tomala T (2014) Secure message transmission on directed networks. *Games Econ Behav* 85:1–18
- Rosenberg D (1998) Duality and Markovian strategies. *Int J Game Theory* 27:577–597
- Rosenberg D, Sorin S (2001) An operator approach to zero-sum repeated games. *Israel J Math* 121:221–246

- Rosenberg D, Solan E, Vieille N (2004) Stochastic games with a single controller and incomplete information. *SIAM J Control Optim* 43:86–110
- Rustichini A (1999) Minimizing regret: the general case. *Games Econom Behav* 29:224–243
- Shalev J (1994) Nonzero-sum two-person repeated games with incomplete information and known-own payoffs. *Games Econom Behav* 7:246–259
- Simon RS (2002) Separation of joint plan equilibrium payoffs from the min-max functions. *Games Econom Behav* 1:79–102
- Simon RS, Spieź S, Toruńczyk H (1995) The existence of equilibria in certain games, separation for families of convex functions and a theorem of Borsuk- Ulam type. *Israel J Math* 92:1–21
- Simon RS, Spieź S, Toruńczyk H (2002) Equilibrium existence and topology in some repeated games with incomplete information. *Trans AMS* 354:5005–5026
- Simon RS, Spieź S, Toruńczyk H (2008) Equilibria in a class of games and topological results implying their existence. *Rev R Acad Cien Serie A Mat* 102:161–179
- Sion M (1958) On general minimax theorems. *Pac J Math* 8:171–176
- Sorin S (1983) Some results on the existence of Nash equilibria for non- zero sum games with incomplete information. *Int J Game Theory* 12:193–205
- Sorin S (1984a) Big match with lack of information on one side (Part I). *Int J Game Theory* 13:201–255
- Sorin S (1984b) On a pair of simultaneous functional equations. *J Math Anal Appl* 98:296–303
- Sorin S (1989) On recursive games without a recursive structure: existence of $\lim_{n \rightarrow \infty}$. *Int J Game Theory* 18:45–55
- Sorin S (1997) Merging, reputation, and repeated games with incomplete information. *Games Econom Behav* 29:274–308
- Sorin S, Zamir S (1985) A 2-person game with lack of information on 1 and 1/2 sides. *Math Oper Res* 10:17–23
- Spinat X (2002) A necessary and sufficient condition for approachability. *Math Oper Res* 27:31–44
- Vieille N (1992) Weak approachability. *Math Oper Res* 17:781–791
- Waternaux C (1983) Solution for a class of repeated games without recursive structure. *Int J Game Theory* 12:129–160
- Zamir S (1971) On the relation between finitely and infinitely repeated games with incomplete information. *Int J Game Theory* 1:179–198
- Zamir S (1973) On repeated games with general information function. *Int J Game Theory* 21:215–229
- Ziliotto B (2016) Zero-sum repeated games: counterexamples to the existence of the asymptotic value and the conjecture $\max \min = \lim \min$. *Ann Probab* 44:1107–1133

Books and Reviews

- Aumann RJ, Maschler M (1995) Repeated games with incomplete information, with the collaboration of R.E. Stearns. M.I.T. Press, 1995 (contains a reedition of chapters of Reports to the U.S. Arms Control and Disarmament Agency ST-80, 116 and 143, *Mathematica*, 1966-1967-1968)
- Forges F (1992) Repeated games of incomplete information: non-zero sum. In: Aumann RJ, Hart S (eds) *Handbook of game theory*, vol I. Elsevier, North-Holland, pp 155–177
- Laraki R, Sorin S (2014) Chapter 2: Advances in zero-sum dynamic games. In: Zamir S, Young P (eds) *Handbook of game theory*, vol IV. Elsevier, North-Holland, pp 27–93
- Laraki R, Renault J, Tomala T (2006) *Théorie des Jeux, Introduction à la théorie des jeux répétées. Editions de l'Ecole Polytechnique, journées X-UPS 2006*. ISBN: 978-2-7302-1366-0, in French (Chapter 3 deals with repeated games with incomplete information)
- Mertens J-F (1987) Repeated games. In: *Proceedings of the international congress of mathematicians*, Berkeley 1986. American Mathematical Society, Dordrecht, pp 1528–1577
- Mertens J-F, Sorin S, Zamir S (1994) Repeated games. CORE discussion paper 9420-9422
- Sorin S (2002) *A first course on zero-sum repeated games. Mathématiques et Applications*. Springer-Verlag Berlin Heidelberg
- Zamir S (1992) Repeated games of incomplete information: zero-sum. In: Aumann RJ, Hart S (eds) *Handbook of game theory*, vol I. Elsevier, North-Holland, pp 109–154

Reputation Effects

George J. Mailath

Department of Economics, University of
Pennsylvania, Philadelphia, USA

Article Outline

Glossary

Definition of the Subject

Introduction

A Canonical Model

Two Long-Lived Players

Future Directions

Bibliography

Glossary

Action type A type of player who is committed to playing a particular action, also called a *commitment type* or *behavioral type*.

Complete information Characteristics of all players are common knowledge.

Flow payoff Stage game payoff.

Imperfect monitoring Past actions of all players are not public information.

Incomplete information Characteristics of some player are not common knowledge.

Long-lived player Player subject to intertemporal incentives, typically has the same horizon as length of the game.

Myopic optimum An action maximizing stage game payoffs.

Nash equilibrium A strategy profile from which no player has a profitable unilateral deviation (i. e., it is self-enforcing).

Nash reversion In a repeated game, permanent play of a stage game Nash equilibrium.

Normalized discounted value The discounted sum of an infinite sequence $\{a_t\}_{t \geq 0}$, calculated as $(1 - \delta) \sum_{t \geq 0} \delta^t a_t$, where $\delta \in (0, 1)$ is the discount value.

Perfect monitoring Past actions of all players are public information.

Repeated game The finite or infinite repetition of a stage game.

Reputation bound The lower bound on equilibrium payoffs of a player that the other player (s) believe may be a simple action type (typically the Stackelberg type).

Short-lived player Player not subject to intertemporal incentives, having a one-period horizon and so is myopically optimizing.

Simple action type An action who plays the same (pure or mixed) stage-game action in every period, regardless of history.

Stackelberg action In a stage game, the action a player would commit to, if that player had the chance to do so, i. e., the optimal commitment action.

Stackelberg type A simple action type that plays the Stackelberg action.

Stage game A game played in one period.

Subgame perfect equilibrium A strategy profile that induces a Nash equilibrium on every subgame of the original game.

Subgame In a repeated game with perfect monitoring, the game following any history.

Type The characteristic of a player that is not common knowledge.

Definition of the Subject

Repeated games have many equilibria, including the repetition of stage game Nash equilibria. At the same time, particularly when monitoring is imperfect, certain plausible outcomes are not consistent with equilibrium. *Reputation effects* is the term used for the impact upon the set of equilibria (typically of a repeated game) of perturbing the game by introducing incomplete information of a particular kind. Specifically, the characteristics of a player are not public information, and the other players believe it is possible that the distinguished player is a type that necessarily plays some action

(typically the Stackelberg action). Reputation effects fall into two classes: “Plausible” phenomena that are not equilibria of the original repeated game are equilibrium phenomena in the presence of incomplete information, and “implausible” equilibria of the original game are not equilibria of the incomplete information game. As such, reputation effects provide an important qualification to the general indeterminacy of equilibria.

Introduction

Repeating play of a stage game often allows for equilibrium behavior inconsistent with equilibrium of that stage game. If the stage game has multiple Nash equilibrium payoffs, a large finite number of repetitions provide sufficient intertemporal incentives for behavior inconsistent with stage-game Nash equilibria to arise in some subgame perfect equilibria. However, many classic games do not have multiple Nash equilibria. For example, mutual defection *DD* is the unique Nash equilibrium of the prisoners’ dilemma, illustrated in Fig. 1.

A standard argument shows that the finitely repeated prisoner’s dilemma has a unique subgame perfect equilibrium, and in this equilibrium, *DD* is played in every period: In any subgame perfect equilibrium, in the last period, *DD* must be played independently of history, since the stage game has a unique Nash equilibrium. Then, since play in the last period is independent of history, there are no intertemporal incentives in the penultimate period, and so *DD* must again be played independently of history. Proceeding recursively, *DD* must be played in every period independently of history. (In fact, the finitely repeated prisoners’ dilemma has a unique Nash equilibrium outcome, given by *DD* in every period.)

This contrasts with intuition, which suggests that if the prisoners’ dilemma were repeated a

sufficiently large (though finite) number of times, the two players would find a way to play cooperatively (*C*) at least in the initial stages. In response, (Kreps et al. 1982) argued that intuition can be rescued in the finitely repeated prisoners’ dilemma by introducing incomplete information. In particular, suppose each player assigns some probability to their opponent being a behavioral type who mechanistically plays tit-for-tat (i. e., plays *C* in the first period or if the opponent had played *C* in the previous period, and plays *D* if the opponent had played *D* in the previous period) rather than being a rational player. No matter how small the probability, if the number of repetitions is large enough, the rational players will play *C* in early periods, and the fraction of periods in which *CC* is played is close to one.

This is the first example of a *reputation effect*: a small degree of incomplete information (of the right kind) both rescues the intuitive *CC* for many periods as an equilibrium outcome, and eliminates the unintuitive always *DD* as one. In the same issue of the *Journal of Economic Theory* containing (Kreps et al. 1982), Kreps and Wilson(1982) and Milgrom and Roberts (1982) explored reputation effects in the finite chain store of Selten (1978), showing that intuition is again rescued, this time by introducing the possibility that the chain store is a “tough” type who always fights entry.

Reputation effects describe the impact upon the set of equilibria of the introduction of small amounts of incomplete information of a particular form into repeated games (and other dynamic games). Reputation effects fall into two classes: “Plausible” phenomena that are not equilibria of the complete information game are equilibrium phenomena in the presence of incomplete information, and “implausible” equilibria of the complete information game are not equilibria of the incomplete information game.

Reputation effects are distinct from the equilibrium phenomenon in complete information repeated games that are sometimes described as capturing *reputations*. In this latter use, an equilibrium of the complete information repeated game is selected, involving actions along the equilibrium path that are not Nash equilibria of the

Reputation Effects,
Fig. 1 The prisoners’ dilemma. The cooperative action is labeled *C*, while defect is labeled *D*

	<i>C</i>	<i>D</i>
<i>C</i>	2, 2	−1, 3
<i>D</i>	3, −1	0, 0

stage game. As usual, incentives to choose these actions are created by attaching less favorable continuation paths to deviations. Players who choose the equilibrium actions are then interpreted as maintaining a reputation for doing so, with a punishment-triggering deviation interpreted as causing the loss of one's reputation. For example, players who cooperate in the *infinitely* repeated prisoners' dilemma are interpreted as having (or maintaining) a cooperative reputation, with any defection destroying that reputation. In this usage, the link between past behavior and expectations of future behavior is an equilibrium phenomenon, holding in some equilibria, but not in others. The notion of reputation is used to interpret an equilibrium strategy profile, but otherwise adds nothing to the formal analysis.

In contrast, the approach underlying reputation effects begins with the assumption that a player is uncertain about key aspects of her opponent. For example, player 2 may not know player 1's payoffs, or may be uncertain about what constraints player 1 faces on his ability to choose various actions. This incomplete information is a device that introduces an intrinsic connection between past behavior and expectations of future behavior. Since incomplete information about players' characteristics can have dramatic effects on the *set* of equilibrium payoffs, reputations in this approach do not describe certain equilibria, but rather place constraints on the set of possible equilibria.

An Example

While reputation effects were first studied in a symmetric example with two long-lived players, they arise in their purest form in infinitely repeated games with one long-lived player playing against a sequence of short-lived players. The chain store game of Selten (1978) is a finitely repeated game in which a chain store (the long-lived player) faces a finite sequence of potential entrants in its different markets. Since each entrant only cares about its own decision, it is short-lived.

Consider the "product-choice" game of Fig. 2. The row player (player 1), who is long-lived, is a firm choosing between high (H) and low (L) effort, while the column player (player 2), who is

Reputation Effects,
Fig. 2 The product-choice game

	h	ℓ
H	2, 3	0, 2
L	3, 0	1, 1

short-lived, is a customer choosing between a high (h) or low (ℓ) priced product. (Mailath and Samuelson (2006) illustrate various aspects of repeated games and reputation effects using this example.) Player 2 prefers the high-priced product if the firm has exerted high effort, but prefers the low-priced product if the firm has not. The firm prefers that customers purchase the high-priced product and is willing to commit to high effort to induce that choice by the customer. In a simultaneous move game, however, the firm cannot observably choose effort before the customer chooses the product. Since high effort is costly, the firm prefers low effort, no matter the choice of the customer.

The stage game has a unique Nash equilibrium, in which the firm exerts low effort and the customer purchases the low-priced product. Suppose the game is played infinitely often, with *perfect monitoring* (i. e., the history of play is public information). The firm is *long-lived* and discounts flow profits by the discount factor $\delta \in (0, 1)$, and is *patient* if δ is close to 1. The role of the customer is taken by a succession of *short-lived* players, each of whom plays the game only once (and so myopically optimizes). It is standard to abuse language by treating the collection of short-lived players as a single myopically optimizing player.

When the firm is sufficiently patient, there is an equilibrium outcome in the repeated game in which the firm always exerts high effort and customers always purchase the high-priced product. The firm is deterred from taking the immediate myopically optimal action of low effort by the prospect of future customers then purchasing the low-priced product. Purchasing the high-priced product is a best response for the customer to high effort, so that no incentive issues arise concerning the customer's behavior. In this equilibrium, the long-lived player's payoff is 2 (the firm's payoffs are calculated as the normalized discounted sum, i. e., as the discounted sum of flow payoffs normalized by $(1 - \delta)$, so that

payoffs in the infinite horizon game are comparable to flow payoffs). However, there are many other equilibria, including one in which low effort is exerted and low price purchased in every period, leading to a payoff of 1 for the long-lived player. Indeed, for $\delta \geq 1/2$, the set of pure-strategy subgame-perfect-equilibrium player 1 payoffs is given by the entire interval (Abreu and Gul 2000; Benabou and Laroque 1992).

Reputation effects effectively rule out any payoff less than 2 as an equilibrium payoff for player 1. Suppose customers are not entirely certain of the characteristics of the firm. More specifically, suppose they attach high probability to the firm's being "normal," that is, having the payoffs given above, but they also entertain some (possibly very small) probability that they face a firm who fortuitously has a technology or some other characteristic that ensures high effort. Refer to the latter as the "*H-action*" type of firm. Since such a type necessarily plays *H* in every period, it is a type described by behavior (not payoffs), and such a type is often called a *behavioral* or *commitment* type.

This is now a game of *incomplete information*, with the customers uncertain of the firm's type. Since the customers assign high probability to the firm being "normal," the game is in some sense close to the game of complete information. None the less, reputation effects are present: For a sufficiently patient firm, in any Nash equilibrium of the repeated game, the firm's payoff cannot be significantly less than 2. This result holds no matter how unlikely customers think the *H-action* type to be, though increasing patience is required from the normal firm as the action type becomes less likely.

The intuition behind this result is most easily seen by considering pure strategy Nash equilibria of the incomplete information game where the customers believe the firm is either the normal or the *H-action* type. In that case, there is no pure strategy Nash equilibrium with a payoff less than 2δ (which is clearly close to 2 for δ close to 1). In the pure strategy Nash equilibrium, either the firm always plays *H*, (in which case, the customers always play *h* and the firm's payoff is 2), or there is a first period (say *t*) in which the firm plays *L*, revealing to future customers that he is

the normal type (since the action type plays *H* in every period). In such an equilibrium, customers play *h* before *t* (since both types of firm are choosing *H*). After observing *H* in period *t*, customers conclude the firm is the *H-action* type. Consequently, as long as *H* is always chosen thereafter, customers subsequently play *h* (since they continue to believe the firm is the *H-action* type, and so necessarily plays *H*). An easy lower bound on the normal firm's equilibrium payoff is then obtained by observing that the normal firm's payoff must be at least the payoff from mimicking the action type in every period. The payoff from such behavior is at least as large as

$$\begin{aligned}
 & \underbrace{(1 - \delta) \sum_{\tau=0}^{t-1} \delta^\tau 2}_{\text{payoff in } \tau < t \text{ from pooling with } H\text{-action type}} \\
 & + \underbrace{(1 - \delta) \delta^t \times 0}_{\text{payoff in } t \text{ from playing } H \text{ when } L \text{ may be myopically optimal}} \\
 & + \underbrace{(1 - \delta) \sum_{\tau=t+1}^{\infty} \delta^\tau 2}_{\text{payoff in } \tau > t \text{ from playing like and being treated as the } H\text{-action type}} \\
 & = (1 - \delta^t)2 + \delta^{t+1}2 = 2 - 2\delta^t(1 - \delta) \\
 & \geq 2 - 2(1 - \delta) = 2\delta.
 \end{aligned}$$

The outcome in which the stage game Nash equilibrium *Lℓ* is played in every period is thus eliminated.

Since reputation effects are motivated by the hypothesis that the short-lived players are uncertain about some aspect of the long-lived player's characteristics, it is important that the results are not sensitive to the precise nature of that uncertainty. In particular, the lower bound on payoffs should not require that the short-lived players *only* assign positive probability to the normal and the *H-action* type (as in the game just analyzed). And it does not: The customers in the example may assign positive probability to the firm being an action type that plays *H* on even periods, and *L* on odd periods, as well as to an action type that plays *H* in every period before some period *t'* (that can depend on

history), and then always plays L . Yet, as long as the customers assign positive probability to the H -action type, for a sufficiently patient firm, in any Nash equilibrium of the repeated game, the firm's payoff cannot be significantly less than 2.

Reputation effects are more powerful in the presence of imperfect monitoring. Suppose that the firm's choice of H or L is not observed by the customers. Instead, the customers observe a public signal $y \in \{y, \bar{y}\}$ at the end of each period, where the signal $\bar{y}$ is realized with probability $p \in (0, 1)$ if the firm chose H , and with the smaller probability $q \in (0, p)$ if the firm chose L . Interpret $\bar{y}$ as a good meal: while customers do not observe effort, they do observe a noisy signal (the quality of the meal) of that effort, with high effort leading to a good meal with higher probability. In the game with complete information, the largest equilibrium payoff to the firm is now given by

$$\bar{v}_1 \equiv 2 - \frac{1-p}{p-q}, \quad (1)$$

reflecting the imperfect monitoring of the firm's actions (the firm is said to be subject to binding moral hazard, see Sect. 7.6 in Mailath and Samuelson (2006)). Since deviations from H cannot be detected for sure, there are no equilibria with the deterministic outcome path of Hh in every period. In some periods after some histories, $L\ell$ must be played in order to provide the appropriate intertemporal incentives to the firm.

As under perfect monitoring, as long as customers assign positive probability to the H -action type in the incomplete information game with imperfect monitoring, for a sufficiently patient firm, in any Nash equilibrium of the repeated game, the firm's payoff cannot be significantly less than 2 (in particular, this lower bound exceeds $\bar{v}_1$). Thus, in this case, reputation effects provide an intuitive lower bound on equilibrium payoffs that both rules out "bad" equilibrium payoffs, as well as rescues outcomes in which Hh occurs in most periods.

Proving that a reputation bound holds in the imperfect monitoring case is considerably more involved than in the perfect monitoring case. In perfect-monitoring games, it is only necessary to analyze the evolution of the customers' beliefs

when always observing H , the action of the H -action type. In contrast, imperfect monitoring requires consideration of belief evolution on all histories that arise with positive probability.

None the less, the intuition is the same: Consider a putative equilibrium in which the normal firm receives a payoff less than $2 - \varepsilon$. Then the normal and action types must be making different choices over the course of the repeated game, since an equilibrium in which they behave identically would induce customers to choose h and would yield a payoff of 2. As in the perfect monitoring case, the normal firm has the option of mimicking the behavior of the H -action type. Suppose the normal firm does so. Since the customers expect the normal type of firm to behave differently from the H -action type, they will more often see signals indicative of the H -action type (rather than the normal type), and so must eventually become convinced that the firm is the H -action type. Hence, in response to this deviation, the customers will eventually play their best response to H of h . While "eventually" may take a while, that time is independent of the equilibrium (indeed of the discount factor), depending only on the imperfection in the monitoring and the prior probability assigned to the H -action type. Then, if the firm is sufficiently patient, the payoff from mimicking the H -action type is arbitrarily close to 2, contradicting the existence of an equilibrium in which the firm's payoff fell short of $2 - \varepsilon$.

At the same time, because monitoring is imperfect, as discussed in section "Temporary Reputation Effects," the reputation effects are necessarily transient. Under general conditions in imperfect-monitoring games, the incomplete information that is at the core of reputation effects is a short-run phenomenon. Player 2 must eventually come to learn player 1's type and continuation play must converge to an equilibrium of the complete information game.

Reputation effects arise for very general specifications of the incomplete information as long as the customers assign strictly positive probability to the H -action type. It is critical, however, that the customers do assign strictly positive probability to the H -action type. For example, in the product-choice game, the set of Nash equilibria of the repeated game is not significantly impacted by

the possibility that the firm is either normal or the L -action type only. While reputation effects per se do not arise from the L -action type, it is still of interest to investigate the impact of such uncertainty on behavior using stronger equilibrium notions, such as Markov perfection (see Mailath and Samuelson (2001)).

A Canonical Model

The Stage Game

The stage game is a two-player simultaneous-move finite game of public monitoring. Player i has action set A_i , $i = 1, 2$. Pure actions for player i are denoted by $a_i \in A_i$, and mixed actions are denoted by $\alpha_i \in \Delta(A_i)$, where $\Delta(A_i)$ is the set of probability distributions over A_i . Player 2's actions are public, while player 1's are potentially private. The public signal of player 1's action, denoted by y is drawn from a finite set Y , with the probability that y is realized under the pure action profile $a \in A \equiv A_1 \times A_2$ denoted by $\rho(y | a)$. Player 1's ex post payoff from the action profile a and signal realization y is $r_1(y, a)$, and so the ex ante (or expected) flow payoff is $u_1(a) \equiv \sum_y r_1(y, a) \rho(y | a)$. Player 2's ex post payoff from the action profile a and signal realization y is $r_2(y, a_2)$, and so the ex ante (or expected) flow payoff is $u_2(a) \equiv \sum_y r_2(y, a_2) \rho(y | a)$. Since player 2's ex post payoff is independent of player 1's actions, player 1's actions only affect player 2's payoffs through the impact on the distribution of the signals and so on ex ante payoffs. While the ex post payoffs r_i play no explicit role in the analysis, they justify the informational assumptions to be made. In particular, the model requires that histories of signals and past actions are the only information players receive, and so it is important that stage game payoffs u_i are not informative about the action choice (and this is the critical feature delivered by the assumptions that ex ante payoffs are not observable and that player 2's ex post payoffs do not depend on a_1).

Perfect monitoring is the special case where $Y = A_1$ and $\rho(y | a) = 1$ if $y = a_1$, and 0 otherwise.

The results in this section hold under significantly weaker monitoring assumptions. In

particular, it is not necessary that the actions of player 2 be public. If these are also imperfectly monitored, then the ex post payoff for player 1 is independent of player 2 actions. Since player 2 is short-lived, when player 2's actions are not public, it is then natural to also assume that the period t player 2 does not know earlier player 2's actions.

The Complete Information Repeated Game

The stage game is infinitely repeated. Player 1 is long-lived, with payoffs given by the normalized discounted value $(1 - \delta) \sum_{t=0}^{\infty} \delta^t u_1^t$, where $\delta \in (0, 1)$ is the discount factor and u_1^t is player 1's period t flow payoff. Player 1 is *patient* if δ is close to 1. As in our example, the role of player 2 is taken by a succession of *short-lived* players, each of whom plays the game only once (and so myopically optimizes).

Player 1's set of private histories is $\mathcal{H}_1 \equiv \bigcup_{t=0}^{\infty} (Y \times A)^t$ and the set of public histories (which coincides with the set of player 2's histories) is $\mathcal{H} \equiv \bigcup_{t=0}^{\infty} (Y \times A_2)^t$. If the game has perfect monitoring, histories $h = (y^0, a^0; y^1, a^1; \dots; y^{t-1}, a^{t-1})$ in which $y \neq a_1^t$ for some $t \leq t - 1$ arise with zero probability, independently of behavior, and so can be ignored. A strategy σ_1 for player 1 specifies a probability distribution over 1's pure action set for each possible private history, i. e., $\sigma_1 : \mathcal{H}_1 \rightarrow \Delta(A_1)$. A strategy σ_2 for player 2 specifies a probability distribution over 2's pure action set for each possible public history, i. e., $\sigma_2 : \mathcal{H} \rightarrow \Delta(A_2)$.

Definition 1 The strategy profile (σ_1^*, σ_2^*) is a *Nash equilibrium* if

1. there does not exist a strategy σ_1 yielding a strictly higher payoff for player 1 when player 2 plays σ_2^* , and
2. in all periods t , after any history $h^t \in \mathcal{H}$ arising with positive probability under (σ_1^*, σ_2^*) , $\sigma_2^*(h^t)$ maximizes $E[u_2(\sigma_1^*(h_1^t), a_1) | h^t]$, where the expectation is taken over the period t -private histories that player 1 may have observed.

The Incomplete Information Repeated Game

In the incomplete information game, the type of player 1 is unknown to player 2. A possible type of player 1 is denoted by $\xi \in \Xi$, where Ξ is a finite

or countable set (see Fudenberg and Levine (1992) for the uncountable case). Player 2's prior belief about 1's type is given by the distribution μ , with support Ξ . The set of types is partitioned into a set of *payoff types* Ξ_1 , and a set of *action types* $\Xi_2 \equiv \Xi \setminus \Xi_1$. Payoff types maximize the average discounted value of payoffs, which depend on their type and which may be nonstationary,

$$u_1 : A_1 \times A_2 \times \Xi_1 \times \mathbb{N}_0 \rightarrow \mathbb{R}.$$

Type $\xi_0 \in \Xi_1$ is the *normal type* of player 1, who happens to have a stationary payoff function, given by the stage game in the benchmark game of complete information,

$$u_1(a, \xi_0, t) = u_1(a) \quad \forall a \in A, \forall t \in \mathbb{N}_0.$$

It is standard to think of the prior probability $\mu(\xi_0)$ as being relatively large, so the games of incomplete information are a seemingly small departure from the underlying game of complete information, though there is no requirement that this be the case.

Action types (also called *commitment* or *behavioral types*) do not have payoffs, and simply play a specified repeated game strategy. For any repeated-game strategy from the complete information game, $\hat{\sigma}_1 : \mathcal{H}_1 \rightarrow \Delta(A_1)$, denote by $\xi(\hat{\sigma}_1)$ the action type committed to the strategy $\hat{\sigma}_1$. In general, a commitment type of player 1 can be committed to any strategy in the repeated game. If the strategy in question plays the same (pure or mixed) stage-game action in every period, regardless of history, that type is called a *simple action type*. For example, the H -action type in the product-choice game is a simple action type. The (simple action) type that plays the pure action a 1 in every period is denoted by $\xi(a_1)$ and similarly the simple action type committed to $\alpha_1 \in \Delta(A_1)$ is denoted by $\xi(\alpha_1)$. As will be seen soon, allowing for mixed action types is an important generalization from simple pure types.

A strategy for player 1, also denoted by $\sigma_1 : \mathcal{H}_1 \times \Xi \rightarrow \Delta(A_1)$, specifies for each type $\xi \in \Xi$ a repeated game strategy such that for all $\xi(\hat{\sigma}_1) \in \Xi_2$, the strategy $\hat{\sigma}_1$ is specified. A strategy σ_2 for player 2 is as in the complete information game, i. e., $\sigma_2 : H \rightarrow \Delta(A_2)$.

Definition 2 The strategy profile (σ_1^*, σ_2^*) is a *Nash equilibrium* of the incomplete information game if

1. for all $\xi \in \Xi_1$, there does not exist a repeated game strategy σ_1 yielding a strictly higher payoff for payoff type ξ of player 1 when player 2 plays σ_2^* , and
2. in all periods t , after any history $h' \in \mathcal{H}$ arising with positive probability under (σ_1^*, σ_2^*) and μ , $\sigma_2^*(h')$ maximizes $E[u_2(\sigma_1^*(h'_1, \xi), a_1) \mid h']$, where the expectation is taken over both the period t -private histories that player 1 may have observed and player 1's type.

Example 1 Consider the product-choice game (Fig. 2) under perfect monitoring. The firm is willing to commit to H to induce h from customers. This incentive to commit is best illustrated by considering a sequential version of the product-choice game: The firm first publicly commits to an effort, and then the customer chooses between h and ℓ , knowing the firm's choice. In this sequential game, the firm chooses H in the unique subgame perfect equilibrium. Since Stackelberg (1934) was the first investigation of such leader-follower interactions, it is traditional to call H the Stackelberg action, and the H -action type of player 1 the Stackelberg type, with associated Stackelberg payoff 2. Suppose $\Xi = \{\xi_0, \xi(H), \xi(L)\}$. For $\delta \geq 1/2$, the grim trigger strategy profile of always playing Hh , with deviations punished by Nash reversion, is a subgame perfect equilibrium of the complete information game. Consider the following adaptation of this profile in the incomplete information game:

$$\begin{aligned} \sigma_1(h', \xi) &= \begin{cases} H, & \text{if } \xi = \xi(H), \text{ or } \xi = \xi_0 \text{ and} \\ & a^\tau = Hh \text{ for all } \tau < t, \\ L, & \text{otherwise,} \end{cases} \\ \text{and } \sigma_2(h') &= \begin{cases} h, & \text{if } a^\tau = Hh \text{ for all } \tau < t, \\ \ell, & \text{otherwise.} \end{cases} \end{aligned}$$

In other words, player 2 and the normal type of player 1 follow the strategies from the Nash-reversion equilibrium in the complete information game, and the action types $\xi(H)$ and $\xi(L)$ play their actions.

This is a Nash equilibrium for $\delta \geq 1/2$ and $\mu(\xi(L)) < 1/2$. The restriction on $\mu(\xi(L))$ ensures that player 2 finds h optimal in period 0. Should player 2 ever observe L , then Bayes' rule causes her to place probability 1 on type $\xi(L)$ (if L is observed in the first period) or the normal type (if L is first played in a subsequent period), making her participation in Nash reversion optimal. The restriction on δ ensures that Nash reversion provides sufficient incentive to make H optimal for the normal player 1. After observing $a_1^0 = H$ in period 0, player 2 assigns zero probability to $\xi = \xi(L)$. However, the posterior probability that 2 assigns to the Stackelberg type does not converge to 1. In period 0, the prior probability is $\mu(\xi(H))$. After one observation of H , the posterior increases to $\mu(\xi^*) / [\mu(\xi^*) + \mu(\xi_0)]$, after which it is constant. By stipulating that an observation of H in a history in which L has previously been observed causes player 2 to place probability one on the normal type of player 1, a specification of player 2's beliefs that is consistent with sequentiality is obtained.

As seen in the introduction, for δ close to 1, $\sigma_1(h^t, \xi_0) = L$ for all h^t is not part of any Nash equilibrium.

The Reputation Bound

Which type would the normal type most like to be treated as? Player 1's *pure-action Stackelberg payoff* is defined as

$$v_1^* = \sup_{a_1 \in A} \min_{a_2 \in B(a_1)} u_1(a_1, a_2), \quad (2)$$

where $B(a_1) = \arg \max_{a_2} u_2(a_1, a_2)$ is the set of player 2 myopic best replies to a . If the supremum is achieved by some action a_1^* , that action is an associated Stackelberg action,

$$a_1^* \in \arg \max_{a_1 \in A_1} \min_{a_2 \in B(a_1)} u_1(a_1, a_2).$$

This is a pure action to which player 1 would commit, if player 1 had the chance to do so (and hence the name "Stackelberg" action, see the discussion in Example 1), given that such a commitment induces a best response from player 2. If there is more than one such action for player 1, the action can be chosen arbitrarily.

However, player 1 would typically prefer to commit to a mixed action. In the product-choice

game, for example, a commitment by player 1 to mixing between H and L , with slightly larger probability on H , still induces player 2 to choose h and gives player 1 a larger payoff than a commitment to H . Define the mixed-action Stackelberg payoff as

$$v_1^{**} \equiv \sup_{\alpha_1 \in \Delta(A_1)} \min_{\alpha_2 \in B(\alpha_1)} u_1(\alpha_1, \alpha_2), \quad (3)$$

where $B(a_1) = \arg \max_{a_2} u_2(a_1, a_2)$ is the set of player 2's best responses to α . In the product choice game, $v_1^* = 2$, while $v_1^{**} = 5/2$. Typically, the supremum is not achieved by any mixed action, and so there is no mixed-action Stackelberg type. However, there are mixed action types that, if player 2 is convinced she is facing such a type, will yield payoffs arbitrarily close to the mixed-action Stackelberg payoff.

As with imperfect monitoring, simple mixed action types under perfect monitoring raise issues of monitoring, since a deviation by the normal type from the distribution α_1 of a mixed action type $\xi(\alpha_1)$, to some action in the support cannot be detected. However, when monitoring of the pure actions is perfect, it is possible to statistically detect deviations, and this will be enough to imply the appropriate reputation lower bound.

When monitoring is imperfect, the public signals are statistically informative about the actions of the long-lived player under the next assumption (Lemma 1).

Assumption 1 For all $a_2 \in A_2$, the collection of probability distributions $\{\rho(y \mid (a_1, a_2)) : a_1 \in A_1\}$ is linearly independent.

This assumption is trivially satisfied in the perfect monitoring case. Reputation effects still exist when this assumption fails, but the bounds are more complicated to calculate (see Fudenberg and Levine 1992 or Sect. 15.4.1 in Mailath and Samuelson 2006).

Fixing an action for player 2, a_2 , the mixed action α_1 implies the signal distribution $\sum_{a_1} \rho(y \mid (a_1, a_2)) \alpha_1(a_1)$.

Lemma 1 Suppose ρ satisfies Assumption 1. Then, if for some a_2 ,

$$\begin{aligned} & \sum_{a_1} \rho(y \mid (a_1, a_2)) \alpha_1(a_1) \\ &= \sum_{a_1} \rho(y \mid (a_1, a_2)) \alpha_1(a_1), \forall y, \end{aligned} \quad (4)$$

then $\alpha_1 = \alpha'_1$.

Proof Suppose 4 holds for some a_2 . Let R denote the $|Y| \times |A_1|$ matrix whose y - a_1 element is given by $\rho(y \mid (a_1, a_2))$ (so that the a_1 -column is the probability distribution on Y implied by the action profile a). Then, 4 can be written as $R\alpha_1 = R\alpha'_1$, or more simply as $R(\alpha_1 - \alpha'_1) = 0$. By Assumption 1, R has full column rank, and so $x = 0$ is the only vector $x \in R^{|A_1|}$ solving $Rx = 0$.

Consequently, if player 2 believes that the long-lived player's behavior implies a distribution over the signals close to the distribution implied by some particular action α'_1 , then player 2 must believe that the long-lived player's action is also close to α'_1 . Since A_2 is finite, this then implies that when player 2 is best responding to some belief about the long-lived player's behavior implying a distribution over signals sufficiently close to the distribution implied by α'_1 , then player 2 is in fact best responding to α'_1 .

We are now in a position to state the main reputation bound result. Let $v_{-1}(\xi_0, \mu, \delta)$ be the infimum over the set of the normal player 1's pay-offs in any (pure or mixed) Nash equilibrium in the incomplete information repeated game, given the distribution μ over types and the discount factor δ .

Proposition 1 (Fudenberg and Levine (1989, 1992)) Suppose ρ satisfies Assumption 1 and let $\hat{\xi}$ denote the simple action type that always plays $\hat{\alpha}_1 \in \Delta(A_1)$. Suppose $\mu(\xi_0), \mu(\hat{\xi}) > 0$. For every $\eta > 0$, there is a value K such that for all δ ,

$$\begin{aligned} v_{-1}(\xi_0, \mu, \delta) &\geq (1 - \eta)\delta^K \min_{\alpha_2 \in B(\hat{\alpha}_1)} u_1(\hat{\alpha}_1, \alpha_2) \\ &\quad + (1 - (1 - \eta)\delta^K) \min_{a \in A} u_1(a). \end{aligned} \quad (5)$$

This immediately yields the pure action Stackelberg reputation bound. Fix $\varepsilon > 0$. Taking $\hat{\alpha}_1$ in the proposition as the degenerate mixture

that plays the Stackelberg action α_1^* with probability 1, Eq. 5 becomes

$$\begin{aligned} v_{-1}(\xi_0, \mu, \delta) &\geq (1 - \eta)\delta^K v_1^* \\ &\quad + (1 - (1 - \eta)\delta^K) \min_{a \in A} u_1(a) \\ &\geq v_1^* - (1 - (1 - \eta)\delta^K) 2M, \end{aligned}$$

where $M \equiv \max_a |u_1(a)|$. This last expression is at least as large as $v_1^* - \varepsilon$ when $\eta < \varepsilon / (2M)$ and δ is sufficiently close to 1. The mixed action Stackelberg reputation bound is also covered:

Corollary 1 Suppose ρ satisfies Assumption 1 and μ assigns positive probability to some sequence of simple types $\{\xi(\alpha_1^k)\}_{k=1}^\infty$ with each α_1^k in $\Delta(A)$ satisfying

$$v_1^{**} = \lim_{k \rightarrow \infty} \min_{\alpha_2 \in B(\alpha_1^k)} u(\alpha_1^k, \alpha_2).$$

For all $\varepsilon' > 0$, there exists $\delta_- < 1$ such that for all $\delta \in (\delta_-, 1)$,

$$v_{-1}(\xi_0, \mu, \delta) \geq v_1^{**} - \varepsilon'.$$

The remainder of this subsection outlines a proof of Proposition 1. Fix a strategy profile (σ_1, σ_2) (which may be Nash, but at this point of the discussion, need not be). The beliefs μ then induce a probability distribution $\mathbf{P}$ on the set of outcomes, which is the set of possible infinite histories (denoted by h^∞) and realized types, $(Y \times A)^\infty \times \Xi \equiv \Omega$. The probability measure $\mathbf{P}$ describes how the short-lived players believe the game will evolve, given their prior beliefs μ about the types of the long-lived player. Let $\hat{\mathbf{P}}$ denote the probability distribution on the set of outcomes induced by (σ_1, σ_2) and the action type $\hat{\xi}$. The probability measure $\hat{\mathbf{P}}$ describes how the short-lived players believe the game will evolve if the long-lived player's type is $\hat{\xi}$. Finally, let $\tilde{\mathbf{P}}$ denote the probability distribution on the set of outcomes induced by (σ_1, σ_2) conditioning on the long-lived player's type not being the action type $\hat{\xi}$. Then, $\mathbf{P} \equiv \hat{\mu} \hat{\mathbf{P}} + (1 - \hat{\mu}) \tilde{\mathbf{P}}$, where $\hat{\mu} \equiv \mu(\hat{\xi})$.

The discussion after Lemma 1 implies that the optimal behavior of the short-lived player in period t is determined by that player's beliefs over the signal realizations in that period. These beliefs

can be viewed as a *one-step ahead prediction* of the signal y that will be realized conditional on the history h_t , $\mathbf{P}(y | h')$. Let $\hat{\mu}'(h') = \mathbf{P}(\hat{\xi} | h')$ denote the posterior probability after observing h_t that the short-lived player assigns to the long-lived player having type $\hat{\xi}$. Note also that if the long-lived player is the action type $\hat{\xi}$, then the true probability of the signal y is $\hat{\mathbf{P}}(y | h') = \rho(y | (H, \sigma_2(h')))$. Then,

$$\begin{aligned} \mathbf{P}(y | h') &= \hat{\mu}'(h')\hat{\mathbf{P}}(y | h') \\ &\quad + (1 - \hat{\mu}'(h'))\tilde{\mathbf{P}}(y | h'). \end{aligned}$$

The key step in the proof of Proposition 1 is a statistical result on merging. The following lemma essentially says that the short-lived players cannot be surprised too many times. Note first that an infinite public history h^∞ can be thought of as a sequence of ever longer finite public histories h_t . Consider the collection of infinite public histories with the property that player 2 often sees histories h_t that lead to very different one-step ahead predictions about the signals under $\tilde{\mathbf{P}}$ and under $\hat{\mathbf{P}}$ and have a “low” posterior that the long-lived player is $\hat{\xi}$. The lemma asserts that if the long-lived player is in fact the action type $\hat{\xi}$, this collection of infinite public histories has low probability. Seeing the signals more likely under $\hat{\xi}$ leads the short-lived players to increase the posterior probability on $\hat{\xi}$. The posterior probability fails to converge to 1 under $\hat{\mathbf{P}}$ only if the play of the types different from $\hat{\xi}$ leads, on average, to a signal distribution similar to that implied by $\hat{\xi}$. For the purely statistical statement and its proof, see Section 15.4.2 in (Mailath and Samuelson 2006).

Lemma 2 For all $\eta, \psi > 0$ and $\mu^\dagger \in (0, 1]$, there exists a positive integer K such that for all $\mu(\hat{\xi}) \in [\mu^\dagger, 1)$, for every strategy $\sigma_1 : \mathcal{H}_1 \times \Xi \rightarrow \Delta(A_1)$ and $\sigma_2 : \mathcal{H} \rightarrow \Delta(A_2)$,

$$\begin{aligned} \hat{\mathbf{P}}\left(h^\infty : \left| \left\{ t \geq 1 : (1 - \hat{\mu}'(h')) \max_y |\tilde{\mathbf{P}}(y | h') - \right. \right. \right. \\ \left. \left. \left. \hat{\mathbf{P}}(y | h') \right| \geq \psi \right\} \geq K \right) \leq \eta. \end{aligned} \quad (6)$$

Note that the bound K holds for all strategy profiles (σ_1, σ_2) and all prior probabilities $\mu(\hat{\xi}) \in [\mu^\dagger, 1)$. This allows us to bound equilibrium payoffs.

Proof of Proposition 1 Fix $\eta > 0$. From Lemma 1, by choosing ψ sufficiently small in Lemma 2, with $\hat{\mathbf{P}}$ -probability at least $1 - \eta$, there are at most K periods in which the short-lived players are not best responding to $\hat{\alpha}_1$.

Since a deviation by the long-lived player to the simple strategy of always playing $\hat{\alpha}_1$ induces the same distribution on public histories as $\hat{\mathbf{P}}$, the long-lived player's expected payoff from such a deviation is bounded below by the right side of 5.

Temporary Reputation Effects

Under perfect monitoring, there are often pooling equilibria in which the normal and some action type of player 1 behave identically on the equilibrium path (as in Example 1). Deviations on the part of the normal player 1 are deterred by the prospect of the resulting punishment. Under imperfect monitoring, such pooling equilibria do not exist. The normal and action types may play identically for a long period of time, but the normal type always eventually has an incentive to cheat at least a little on the commitment strategy, contradicting player 2's belief that player 1 will exhibit commitment behavior. Player 2 must then eventually learn player 1's type.

In addition to Assumption 1, disappearing reputation effects require full support monitoring.

Assumption 2 For all $a \in A, y \in Y, \rho(y | a) > 0$.

This assumption implies that Bayes' rule determines the beliefs of player 2 about the type of player 1 after *all* histories.

Suppose there are only two types of player 1, the normal type ξ_0 and a simple action type $\hat{\xi}$, where $\hat{\xi} = \xi(\hat{\alpha}_1)$ for some $\hat{\alpha}_1 \in \Delta(A_1)$. The analysis is extended to many commitment types in Section 6.1 in Cripps et al. (2004). It is convenient to denote a strategy for player 1 as a pair of functions $\tilde{\sigma}_1$ and $\hat{\sigma}_1$ (so $\tilde{\sigma}_1(h_1) = \hat{\alpha}_1$ for all $h_1 \in \mathcal{H}_1$), the former for the normal type and the latter for the action type.

Recall that $\mathbf{P} \in \Delta(\Omega)$ is the unconditional probability measure induced by the prior μ , and the strategy profile $(\hat{\sigma}_1, \tilde{\sigma}_1, \sigma_2)$, while $\hat{\mathbf{P}}$ is the measure induced by conditioning on $\hat{\xi}$. Since $\{\xi_0\} = \Xi \setminus \{\hat{\xi}\}$, $\tilde{\mathbf{P}}$ is the measure induced by

conditioning on ξ_0 . That is, $\hat{\mathbf{P}}$ is induced by the strategy profile $\hat{\sigma} = (\hat{\sigma}_1, \sigma_2)$ and $\hat{\mathbf{P}}$ by $\tilde{\sigma} = (\tilde{\sigma}_1, \sigma_2)$, describing how play evolves when player 1 is the commitment and normal type, respectively.

The action of the commitment type satisfies the following assumption.

Assumption 3 Player 2 has a unique stage-game best response to $\hat{\alpha}_1$ (denoted by $\hat{\alpha}_2$) and $\hat{\alpha} \equiv (\hat{\alpha}_1, \hat{\alpha}_2)$ is not a stage-game Nash equilibrium.

Let $\hat{\sigma}_2$ denote the strategy of playing the unique best response $\hat{\alpha}_2$ to $\hat{\alpha}_1$ in each period independently of history. Since $\hat{\alpha}$ is not a stage-game Nash equilibrium, $(\hat{\sigma}_1, \hat{\sigma}_2)$ is not a Nash equilibrium of the complete information infinite horizon game.

Proposition 2 ((Cripps et al. 2004)) Suppose the monitoring distribution ρ satisfies Assumptions 1 and 2, and the commitment action $\hat{\alpha}_1$ satisfies Assumption 3. In any Nash equilibrium of the game with incomplete information, the posterior probability assigned by player 2 to the commitment type, $\hat{\mu}^t$, converges to zero under $\tilde{\mathbf{P}}$, i. e.,

$$\hat{\mu}^t(h^t) \rightarrow 0, \quad \tilde{\mathbf{P}} - \text{a.s.}$$

The intuition is straightforward: Suppose there is a Nash equilibrium of the incomplete information game in which both the normal and the action type receive positive probability in the limit (on a positive probability set of histories). On this set of histories, player 2 cannot distinguish between signals generated by the two types (otherwise player 2 could ascertain which type she is facing), and hence must believe that the normal and action types are playing the same strategies on average. But then player 2 must play a best response to this strategy, and hence to the action type. Since the action type's behavior is not a best response for the normal type (to this player 2 behavior), player 1 must eventually find it optimal to *not* play the action-type strategy, contradicting player 2's beliefs.

Assumption 3 requires a unique best response to $\hat{\alpha}_1$. For example, in the product-choice game, every action for player 2 is a best response to player 1's mixture α'_1 that assigns equal

probability to H and L . This indifference can be exploited to construct an equilibrium in which (the normal) player 1 plays α'_1 after every history (Section 7.6.2 in (Mailath and Samuelson 2006)). This will still be an equilibrium in the game of incomplete information in which the commitment type plays α'_1 , with the identical play of the normal and commitment types ensuring that player 2 never learns player 1's type. In contrast, player 2 has a unique best response to any other mixture on the part of player 1. Therefore, if the commitment type is committed to any mixed action other than α'_1 , player 2 will eventually learn player 1's type.

As in Proposition 1, a key step in the proof of Proposition 2 is a purely statistical result on updating. Either player 2's expectation (given her history) of the strategy played by the normal type $\tilde{E}[\tilde{\sigma}_1^t | h^t]$, where $\tilde{E}$ denotes expectation with respect to $\tilde{\mathbf{P}}$ is in the limit identical to the strategy played by the action type ($\hat{\alpha}_1$), or player 2's posterior probability that player 1 is the action type ($\hat{\mu}(h^t)$) converges to zero (given that player 1 is indeed normal). This is a merging argument and closely related to Lemma 2. If the distributions generating player 2's signals are different for the normal and action type, then these signals provide information that player 2 will use in updating her posterior beliefs about the type she faces. This (converging, since beliefs are a martingale) belief can converge to an interior probability only if the distributions generating the signals are asymptotically uninformative, which requires that they be asymptotically identical.

Lemma 3 Suppose the monitoring distribution ρ satisfies Assumptions 1 and 2. Then in any Nash equilibrium,

$$\begin{aligned} \lim_{x \rightarrow \infty} \hat{\mu}^t \max_{a_1} |\hat{\alpha}(a) - \tilde{E}[\tilde{\sigma}_1^t(a_1) | h^t]| \\ = 0, \tilde{\mathbf{P}} - \text{a.s.} \end{aligned} \quad (7)$$

Given Proposition 2, it should be expected that continuation play converges to an equilibrium of the complete information game, and this is indeed the case. See Theorem 2 (Cripps et al. 2004) for the formal statement.

Proposition 2 leaves open the possibility that for any period T , there may be equilibria in which uncertainty about player 1's type survives beyond T , even though such uncertainty asymptotically disappears in any equilibrium. This possibility cannot arise. The existence of a sequence of Nash equilibria with uncertainty about player 1's type persisting beyond period $T \rightarrow \infty$ would imply the (contradictory) existence of a limiting Nash equilibrium in which uncertainty about player 1's type persists.

Proposition 3 ((Cripps et al. 2007)) Suppose the monitoring distribution ρ satisfies Assumptions 1 and 2, and the commitment action $\hat{\alpha}_1$ satisfies Assumption 3. For all $\varepsilon > 0$, there exists T such that for any Nash equilibrium of the game with incomplete information,

$$\tilde{\mathbf{P}}(\hat{\mu}^t < \varepsilon, \forall t > T) > 1 - \varepsilon.$$

Example 2 Recall that in the product-choice game, the unique player 2 best response to H is to play h , and Hh is not a stage-game Nash equilibrium. Proposition 1 ensures that the normal player 1's expected value in the repeated game of incomplete information with the H -action type is arbitrarily close to 2, when player 1 is very patient. In particular, if the normal player 1 plays H in every period, then player 2 will at least eventually play her best response of h . If the normal player 1 persisted in mimicking the action type by playing H in each period, this behavior would persist indefinitely. It is the feasibility of such a strategy that lies at the heart of the reputation bounds on expected payoffs. However, this strategy is not optimal. Instead, player 1 does even better by attaching some probability to L , occasionally reaping the rewards of his reputation by earning a stage-game payoff even larger than 2. The result of such equilibrium behavior, however, is that player 2 must eventually learn player 1's type. The *continuation* payoff is then bounded below 2 (recall (1)).

Reputation effects arise when player 2 is uncertain about player 1's type, and there may well be a long period of time during which player 2 is sufficiently uncertain of player 1's type (relative to the discount factor), and in which play does not resemble an equilibrium of the complete information

game. Eventually, however, such behavior must give way to a regime in which player 2 is (correctly) convinced of player 1's type.

For any prior probability $\hat{\mu}$ that the long-lived player is the commitment type and for any $\varepsilon > 0$, there is a discount factor δ sufficiently large that player 1's expected payoff is close to the commitment-type payoff. This holds no matter how small $\hat{\mu}$. However, for any fixed δ and in any equilibrium, there is a time at which the posterior probability attached to the commitment type has dropped below the corresponding critical value of $\hat{\mu}$, becoming too small (relative to δ) for reputation effects to operate.

A reasonable response to the results on disappearing reputation effects is that a model of *long-run* reputations should incorporate some mechanism by which the uncertainty about types is continually replenished. For example, Holmström (1982), (Cole et al. 1995), Mailath and Samuelson (2001), and Phelan (2006) assume that the type of the long-lived player is governed by a stochastic process rather than being determined once and for all at the beginning of the game. In such a situation, reputation effects can indeed have long-run implications.

Reputation as a State

The posterior probability that short-lived players assign to player 1 being $\hat{\xi}$ is sometimes interpreted as player 1's reputation, particularly if $\hat{\xi}$ is the Stackelberg type. When Ξ contains only the normal type and $\hat{\xi}$, the posterior belief $\hat{\mu}^t$ is a state variable of the game, and attention is sometimes restricted to Markov strategies (i. e., strategies that only depend on histories through their impact on the posterior beliefs of the short-lived players). An informative example is Benabou and Laroque (1992), who study the Markov perfect equilibria of a game in which the uninformed players respond continuously to their beliefs. They show that the informed player eventually reveals his type in any Markov perfect equilibrium. On the other hand, Markov equilibria need not exist in finitely repeated reputation games (Section 17.3 in (Mailath and Samuelson 2006)).

The literature on reputation effects has typically not restricted attention to Markov strategies, since the results do not require the restriction.

Two Long-Lived Players

The introduction of nontrivial intertemporal incentives for the uninformed player significantly reduces reputation effects. For example, when only simple Stackelberg types are considered, the Stackelberg payoff may not bound equilibrium payoffs. The situation is further complicated by the possibility of non-simple commitment types (i. e., types that follow nonstationary strategies).

Consider applying the logic from section “[The Reputation Bound](#)” to obtain the Stackelberg reputation bound when both players are long-lived and player 1’s characteristics are unknown, under perfect monitoring. The first step is to demonstrate that, if the normal player 1 persistently plays the Stackelberg action and there exists a type committed to that action, then player 2 must eventually attach high probability to the event that the Stackelberg action is played in the future. This argument, a simple version of Lemma 2, depends only upon the properties of Bayesian belief revision, independently of whether the person holding the beliefs is a long-lived or short-lived player.

When player 2 is short-lived, the next step is to note that if she expects the Stackelberg action, then she will play a best response to this action. If player 2 is instead a long-lived player, she may have an incentive to play something other than a best response to the Stackelberg type.

The key step when working with two long-lived players is thus to establish conditions under which, as player 2 becomes increasingly convinced that the Stackelberg action will appear, player 2 must eventually play a best response to that action. One might begin such an argument by observing that, as long as player 2 discounts, any losses from not playing a current best response must be recouped within a finite length of time. But if player 2 is “very” convinced that the Stackelberg action will be played not only now but for sufficiently many periods to come, there will be no opportunity to accumulate subsequent gains, and hence player 2 might just as well play a stage-game best response.

Once it is shown that player 2 is best responding to the Stackelberg action, the remainder of the argument proceeds as in the case of a short-lived player 2. The normal player 1 must eventually receive very nearly the Stackelberg payoff in each

period of the repeated game. By making player 1 sufficiently patient (*relative to player 2*, so that discount factors differ), this consideration dominates player 1’s payoffs, putting a lower bound on the latter. Hence, the obvious handling of discount factors is to fix player 2’s discount factor δ_2 , and to consider the limit as player 1 becomes patient, i. e., δ_1 approaching one.

This intuition misses the following possibility. Player 2 may be choosing something other than a best response to the Stackelberg action out of fear that a current best response may trigger a disastrous future punishment. This punishment would not appear if player 2 faced the Stackelberg type, but player 2 can be made confident only that she faces the Stackelberg *action*, not the Stackelberg type. The fact that the punishment lies off the equilibrium path makes it difficult to assuage player 2’s fear of such punishments. Short-lived players in the same situation are similarly uncertain about the future ramifications of best responding, but being short-lived, this uncertainty does not affect their behavior.

Consequently, reputation effects are typically weak with two long-lived players under perfect monitoring: (Celentani et al. 1996) and Cripps and Thomas (1997), describe examples with only the normal and the Stackelberg types of player 1, in which the future play of the normal player 1 is used to punish player 2 for choosing a best response to the Stackelberg action when she is not supposed to, and player 1’s payoff is significantly below the Stackelberg payoff. Moreover, the robustness of reputation effects to additional types beyond the Stackelberg type, a crucial feature of settings with one long-lived player, does not hold with two long-lived players. Schmidt (1993b) showed that the possibility of a “punishment” type can prevent player 2 best responding to the Stackelberg action, while Evans and Thomas (1997) showed that the Stackelberg bound is valid if in addition to the Stackelberg type, there is an action type who punishes player 2 for *not* behaving appropriately (see Sections 16.1 and 16.5 in (Mailath and Samuelson 2006)).

Imperfect monitoring (of *both* players’ actions), on the other hand, rescues reputation effects. With a sufficiently rich set of commitment types, player 1 can be assured of at least his Stackelberg payoff.

Indeed, player 1 can often be assured of an even higher payoff, in the presence of commitment types who play nonstationary strategies (Celentani et al. 1996). At the same time, these reputation effects are temporary (Theorem 2 in (Cripps et al. 2007)).

Finally, there is a literature on reputation effects in bargaining games (see (Abreu and Gul 2000; Chatterjee and Samuelson 1987, 1988; Schmidt 1993a)), where the issues described above are further complicated by the need to deal with the bargaining model itself.

Future Directions

The detailed structure of equilibria of the incomplete information game is not well understood, even for the canonical game of section “[A Canonical Model](#).” A more complete description of the structure of equilibria is needed.

While much of the discussion was phrased in terms of the Stackelberg type, Proposition 1 provides a reputation bound for *any* action type. While in some settings, it is natural that the uninformed players assign strictly positive probability to the Stackelberg type, it is not natural in other settings. A model endogenizing the nature of action types would be an important addition to the reputation literature.

Finally, while the results on reputation effects with two long-lived players are discouraging, there is still the possibility that some modification of the model will rescue reputation effects in this important setting.

Acknowledgments I thank Eduardo Faingold, KyungMin Kim, Antonio Penta, and Larry Samuelson for helpful comments.

Bibliography

- Abreu D, Gul F (2000) Bargaining and reputation. *Econometrica* 68(1):85–117
- Benabou R, Laroque G (1992) Using privileged information to manipulate markets: insiders, gurus, and credibility. *Q J Econ* 107(3):921–958
- Celentani M, Fudenberg D, Levine DK, Pesendorfer W (1996) Maintaining a reputation against a long-lived opponent. *Econometrica* 64(3):691–704
- Chatterjee K, Samuelson L (1987) Bargaining with two-sided incomplete information: an infinite horizon model with alternating offers. *Rev Econ Stud* 54(2):175–192
- Chatterjee K, Samuelson L (1988) Bargaining with two-sided incomplete information: the unrestricted offers case. *Oper Res* 36(4):605–638
- Cole HL, Dow J, English WB (1995) Default, settlement, and signalling: lending resumption in a reputational model of sovereign debt. *Int Econ Rev* 36(2):365–385
- Cripps MW, Thomas JP (1997) Reputation and perfection in repeated common interest games. *Games Econ Behav* 18(2):141–158
- Cripps MW, Mailath GJ, Samuelson L (2004) Imperfect monitoring and impermanent reputations. *Econometrica* 72(2):407–432
- Cripps MW, Mailath GJ, Samuelson L (2007) Disappearing private reputations in long-run relationships. *J Econ Theory* 134(1):287–316
- Evans R, Thomas JP (1997) Reputation and experimentation in repeated games with two long-run players. *Econometrica* 65(5):1153–1173
- Fudenberg D, Levine DK (1989) Reputation and equilibrium selection in games with a patient player. *Econometrica* 57(4):759–778
- Fudenberg D, Levine DK (1992) Maintaining a reputation when strategies are imperfectly observed. *Rev Econ Stud* 59(3):561–579
- Holmström B (1982) Managerial incentive problems: a dynamic perspective. In: *Essays in economics and management in honour of Lars Wahlbeck*. Swedish School of Economics and Business Administration, Helsinki, pp 209–230. Published in: *Rev Econ Stud* 66(1):169–182
- Kreps D, Wilson R (1982) Reputation and imperfect information. *J Econ Theory* 27:253–279
- Kreps D, Milgrom PR, Roberts DJ, Wilson R (1982) Rational cooperation in the finitely repeated prisoner’s dilemma. *J Econ Theory* 27:245–252
- Mailath GJ, Samuelson L (2001) Who wants a good reputation? *Rev Econ Stud* 68(2):415–441
- Mailath GJ, Samuelson L (2006) *Repeated games and reputations: long-run relationships*. Oxford University Press, New York
- Milgrom PR, Roberts DJ (1982) Limit pricing and entry under incomplete information: an equilibrium analysis. *Econometrica* 50:443–459
- Phelan C (2006) Public trust and government betrayal. *J Econ Theory* 130(1):27–43
- Schmidt KM (1993a) Commitment through incomplete information in a simple repeated bargaining game. *J Econ Theory* 60(1):114–139
- Schmidt KM (1993b) Reputation and equilibrium characterization in repeated games of conflicting interests. *Econometrica* 61(2):325–351
- Selten R (1978) Chain-store paradox. *Theory Decis* 9:127–159
- Stackelberg HV (1934) *Marktform und Gleichgewicht*. Springer, Vienna

Zero-Sum Two Person Games

T. E. S. Raghavan
Department of Mathematics, Statistics and
Computer Science, University of Illinois,
Chicago, IL, USA

Article Outline

Introduction
Games with Perfect Information
The Game of Hex
Approximate Fixed Points
An Application of the Algorithm
Extensive Games and Normal Form Reduction
Saddle Point
Mixed Strategy and Minimax Theorem
Historical Remarks
Solving for Value and Optimal Strategies via
Linear Programming
Simplex Algorithm
Fictitious Play
Search Games
Search Games on Trees
Umbrella Folding Algorithm
Completely Mixed Games and Perron's Theorem
on Positive Matrices
Behavior Strategies in Games with Perfect Recall
Efficient Computation of Behavior Strategies
General Minimax Theorems
Geometric Consequences
Ky Fan-Sion Minimax Theorems
Applications of Infinite Games
General Minimax Theorem and Statistical
Estimation
Borel's Poker Model
War Duels and Discontinuous Payoffs on the Unit
Square
Bibliography

Introduction

Conflicts are an inevitable part of human existence. This is a consequence of the competitive stances of greed and the scarcity of resources, which are rarely balanced without open conflict. Epic poems of the Greek, Roman, and Indian civilizations which document wars between nation-states or clans reinforce the historical legitimacy of this statement. It can be deduced that domination is the recurring theme in human conflicts. In a primitive sense this is historically observed in the domination of men over women across cultures while on a more refined level it can be observed in the imperialistic ambitions of nation-state actors. In modern times, a new source of conflict has emerged on an international scale in the form of economic competition between multinational corporations.

While conflicts will continue to be a perennial part of human existence, the real question at hand is how to formalize mathematically such conflicts in order to have a grip on potential solutions. We can use mock conflicts in the form of parlor games to understand and evaluate solutions for real conflicts. Conflicts are unresolvable when the participants have no say in the course of action. For example one can lose interest in a parlor game whose entire course of action is dictated by chance. Examples of such games are Chutes and Ladders, Trade, Trouble etc. Quite a few parlor games combine tactical decisions with chance moves. The game Le Her and the game of Parcheesi are typical examples. An outstanding example in this category is the game of backgammon, a remarkably deep game. In chess, the player who moves first is usually determined by a coin toss, but the rest of the game is determined entirely by the decisions of the two players. In such games, players make strategic decisions and attempt to gain an advantage over their opponents.

A game played by two rational players is called zero-sum if one player's gain is the other player's loss. Chess, Checkers, Gin Rummy, Two-finger Morra, and Tic-Tac-Toe are all examples of zero-sum two-person games. Business competition between two major airlines, two major publishers, or two major automobile manufacturers can be modeled as a zero-sum two-person games (even if the outcome is not precisely zero-sum). Zero-sum games can be used to construct Nash equilibria in many dynamic non-zero-sum games (Thuijsman and Raghavan 1997).

Games with Perfect Information

Emptying a Box

Example 1 A box contains 15 pebbles. Players I and II remove between one and four pebbles from the box in alternating turns. Player I goes first, and the game ends when all pebbles have been removed. The player who empties the box on his turn is the winner, and he receives \$1 from his opponent.

The players can decide in advance how many pebbles to remove in each of their turn. Suppose a player finds x pebbles in the box when it is his turn. He can decide to remove 1, 2, 3 or at most 4 pebbles. Thus a *strategy* for a player is any function f whose domain is $X = \{1, 2, \dots, 15\}$ and range is $R \subseteq \{1, 2, 3, 4\}$ such that $f(x) \leq \min(x, 4)$. Given strategies f, g for players I and II respectively, the game evolves by executing the strategies decided in advance. For example if, say

$$f(x) = \begin{cases} 2 & \text{if } x \text{ is even} \\ 1 & \text{if } x \text{ is odd,} \end{cases}$$

$$g(x) = \begin{cases} 3 & \text{if } x \geq 3 \\ x & \text{otherwise.} \end{cases}$$

The alternate depletions lead to the following scenario

move by	I	II	I	II	I	II	I	II
removes	1	3	1	3	1	3	1	2
leaving	14	11	10	7	6	3	2	0

In this case the winner is Player II. Actually in his first move Player II made a bad move by

removing 3 out of 14. Player I could have exploited this. But he did not! Though he made a good second move, he reverted back to his naive strategy and made a bad third move. The question is: Can Player II ensure victory for himself by intelligently choosing a suitable strategy? Indeed Player II can win the game with any strategy satisfying the conditions of g^* where

$$g^*(x) = \begin{cases} 1 & \text{if } x - 1 \text{ is a multiple of 5} \\ 2 & \text{if } x - 2 \text{ is a multiple of 5} \\ 3 & \text{if } x - 3 \text{ is a multiple of 5} \\ 4 & \text{if } x - 4 \text{ is a multiple of 5} \end{cases}$$

Since the game starts with 15 pebbles, Player I must leave either 14 or 13 or 12, or 11 pebbles. Then Player II can in his turn remove 1 or 2 or 3 or 4 pebbles so that the number of pebbles Player I finds is a multiple of 5 at the beginning of his turn. Thus Player II can leave the box empty in the last round and win the game.

Many other combinatorial games could be studied for optimal strategic behavior. We give one more example of a combinatorial game, called the *game of Nim* (Bouton 1902).

Nim Game

Example 2 Three baskets contain 10, 11, and 16 oranges respectively. In alternating turns, Players I and II choose a non-empty basket and remove at least one orange from it. The player may remove as many oranges as he wishes from the chosen basket, up to the number the basket contains. The game ends when the last orange is removed from the last non-empty basket. The player who takes the last orange is the winner.

In this game as in the previous example at any stage the players are fully aware of what has happened so far and what moves have been made. The full history and the state of the game at any instance are known to both players. Such a game is called a game with *perfect information*. How to plan for future moves to one's advantage is not at all clear in this case. Bouton (1902) proposed an ingenious solution to this problem which predates the development of formal game theory.

His solution hinges on the binary representation of any number and the inequality that $1 + 2 +$

$4 + \dots + 2^n < 2^{n+1}$. The numbers 10, 11, 16 have the binary representation

Number	Binary representation
10	= 1010
11	= 1011
16	= 10000
Column totals (in base 10 digits)	= 12021.

Bouton made the following key observations:

1. If at least one column total is an odd number, then the player who is about to make a move can choose one basket and by removing *a suitable number* of oranges leave all column totals even.
2. If at least one basket is nonempty and if all column totals are even, then the player who has to make a move will end up leaving an odd column total.

By looking for the first odd column total from the left, we notice that the basket with 16 oranges is the right choice for Player I. He can remove the left most 1 in the binary expansion of 16 and change all the other binary digits to the right by 0 or 1. The key observation is that the new number is strictly less than the original number. In Player I's move, at least one orange will be removed from a basket. Furthermore, the new column totals can all be made even. If an original column total is even we leave it as it is. If an original column total is odd, we make it even by making any 1 a 0 and any 0 a 1 in those cases which correspond to the basket with 16 oranges. For example the new binary expansion corresponds to removing all but 1 orange from basket 3. We have

10	=	1010
11	=	1011
1	=	0001
Column totals in base 10	=	2022.

In the next move, no matter what a player does, he has to leave one of the 1's a 0 and the column total in that column will be odd and the move is for Player I. Thus Player I will be the first to empty the baskets and win the game.

For the game of Nim we found a constructive and explicit strategy for the winner regardless of any action by the opponent. Sometimes one may be able to assert who should be the winner without knowing any winning strategy for the player!

Definition 3 A zero-sum two person game has perfect information if, at each move, both players know the complete history so far.

There are many variations of nim games and other combinatorial games like Chess and Go that exploit the combinatorial structure of the game or the end games to develop winning strategies. The classic monographs on combinatorial game theory is by Berlekamp et al. (1982) on Winning Ways for your Mathematical Plays, whose mathematical foundations were provided by Conway's earlier book On Numbers and Games. These are often characterized by sequential moves by two players and the outcome is either a win or lose kind. Since the entire history of past moves is common knowledge, the main thrust is in developing winning strategies for such games.

Definition 4 A zero-sum two person game is called a win-lose game if there are no chance moves and the final outcome is either Player I wins or loses (Player II wins) the game. (In other words, there is no way for the game to end in a tie.)

The following is a fundamental theorem of Zermelo (1913).

Theorem 5 Any zero-sum two person perfect information win-lose game Γ with finitely many moves and finitely many choices in each move has a winner with an optimal winning strategy.

Proof Let Player I make the first move. Then depending on the choices available, the game evolves to a new set of subgames which on their own right are also win-lose games of perfect information. Among these subgames the one with the longest play will have fewer moves than the original game. By an induction on the length of the longest play, we can find a winner with a winning strategy, one for each subgame. Each player can develop good strategies for the original

game as follows. Suppose the subgames are $\Gamma_1, \Gamma_2, \dots, \Gamma_k$. Now among these subgames, let Γ_s be a game where Player I can ensure a victory for himself, no matter what Player II does in the subgame. In this case, Player I can determine at the very beginning, the right choice of action which leads to the subgame Γ_s . A good strategy for Player I is simply the choice s in the first move followed by his good strategy in the subgame Γ_s . Player II's strategy for the original game is simply a k -tuple of strategies, one for each subgame. Player II must be ready to use an optimal strategy for the subgame Γ_r in case the first move of Player I leads to playing Γ_r , which is favorable to Player II. Suppose no subgame Γ_s has a winning strategy for Player I. Then Player II will be the winner in each subgame. To achieve this, Player II must use his winning strategy in each subgame they are lead to. Such a k -tuple of winning strategies, one for each subgame, is a winning strategy for Player II for the original game Γ .

The Game of Hex

An interesting win-lose game was made popular by John Nash in late forties among Princeton graduate students. While the original game is aesthetically pleasing with its hexagonal tiles forming a 10×10 rhombus, it is more convenient to use the following equivalent formulation for its mathematical simplicity. The version below is extendable to multi person games and is useful for developing important algorithms (Gale 1979).

Let B_n be a square board consisting of lattice points $\{(i, j): 1 \leq i \leq n, 1 \leq j \leq n\}$. The game involves occupation of unoccupied vertices by players I and II. The board is enlarged with a frame on all sides. The frame $\mathbf{F}$ consists of lattice points $\mathbf{F} = \{(i, j): 0 \leq i \leq n+1, 0 \leq j \leq n+1 \text{ where either } i = 0 \text{ or } n+1, \text{ or } j = 0 \text{ or } n+1\}$. The frame on the west side $W = \{(i, j): i = 0\} \cap \mathbf{F}$ and the frame on the east side $E = \{(i, j): i = n+1\} \cap \mathbf{F}$ are reserved for Player I. Similarly the frame on the south side $S = \{(i, j): 0 < i < n+1, j = 0\} \cap \mathbf{F}$ and frame on the north side $N = \{(i, j): (0 < i < n+1, j = n+1)\} \cap \mathbf{F}$ are reserved for Player II. Two lattice points $P = (x_1, y_1), Q = (x_2, y_2)$ are called *adjacent vertices*


iff either $x_1 \leq x_2, y_1 \leq y_2$ or $x_1 \geq x_2, y_1 \geq y_2$ and $\max(|x_1 - x_2|, |y_1 - y_2|) = 1$. For example the lattice points (4, 10) and (5, 11) are adjacent while (4, 10) and (5, 9) are not adjacent. Six vertices are adjacent to any interior lattice point of the Hex board B_n while lattice points on the frame will have fewer than six adjacent vertices.

The game is played as follows: Players I and II, in alternate turns, choose a vertex from the available set of unoccupied vertices. The aim of Player I is to occupy a bridge of adjacent vertices that links a vertex on the west boundary with a vertex on the east boundary. Player II has a similar objective to connect the north and south boundary with a bridge.

Theorem 6 The game of Hex can never end in a draw. For any $T \subseteq B_n$ occupied by Player I and the complement T^c occupied by Player II, either T contains a winning bridge for Player I or T^c contains a winning bridge for Player II. Further only one can have a winning bridge.

Proof We label any vertex with 1 or 2 depending on who (Player I or Player II) occupies the vertex. Consider triangles Δ formed by vertices that are mutually adjacent to each other. Two such triangles are called *mates* if they share a common side. Either all the 3 vertices of the triangle are occupied by one player or two vertices by one player and the third by the other player. For example if $P = (x_1, y_1), Q = (x_2, y_2), R = (x_3, y_3)$ are adjacent to each other, and if P, Q, R are occupied by say, I, II, and I, they get the labels 1, 2 and 1 respectively. The triangle has exactly 2 sides (PQ and QR) with vertices labeled 1 and 2. The algorithm described below involves entering such a triangle via one side with vertex labels 1 and 2 and exiting via the other side with vertex labels 1 and 2. Suppose we start at the south west corner triangle Δ_0 (in the above figure) with vertex $A = (0, 0)$ occupied by player I (labeled 1), $B = (1, 0)$ occupied by player II (labeled 2), and suppose $C = (1, 1)$ is occupied by, player I (labeled 1). Since we want to stay inside the framed Hex board, the only way to exit Δ_0 via a side with vertices labeled 1 and 2 is to exit via BC . We move to the unique mate triangle Δ_1

Zero-Sum Two Person Games, Fig. 1 Hex path via mate triangles



which shares the common side BC which has vertex labels 1 and 2. The mate triangle Δ_1 has vertices $(1, 0)$, $(1, 1)$, and $(1, 0) + (1, 1) = (2, 1)$. Suppose $D = (2, 1)$ is labeled 2, then we exit via the side CD to the mate triangle with vertices C , D , and $E = (1, 1) - (1, 0) + (2, 1) = (2, 2)$. Each time we find the player of the new vertex with his label, we drop out the other vertex of the same player from the current triangle and move into the new mate triangle. In each iteration there is exactly one new mate triangle to move into. Since in the initial step we had a unique mate triangle to move into from Δ_0 , there is no way for the algorithm to reenter a mate triangle visited earlier. This process must terminate at a vertex on the North or East boundary. One side of these triangles will all have the same label forming a bridge which joins the appropriate boundaries and forms a winning path. The winning player's bridge will obstruct the bridge the losing player attempted to complete. The game of Hex and its winning strategy is a powerful tool in developing algorithms for computing approximate fixed points. Hex is an example of a game where we do know that the

first player can win, but we don't know how (for a sufficiently large board).

Approximate Fixed Points

Let I^2 be the unit square $0 \leq x, y \leq 1$. Given any continuous function: $\mathbf{f} = (f_1, f_2): I^2 \rightarrow I^2$, Brouwer's fixed point theorem asserts the existence of a point (x^*, y^*) such that $\mathbf{f}(x^*, y^*) = (x^*, y^*)$. Our Hex path building algorithm due to Gale (1979) gives a constructive approach to locating an approximate fixed point.

Given $\epsilon > 0$, by uniform continuity we can find a $\delta > \frac{1}{n} > 0$ such that if (i, j) and (i', j') are adjacent vertices of a Hex board B_n , then

$$\begin{aligned} \left| f_1\left(\frac{i}{n}, \frac{j}{n}\right) - f_1\left(\frac{i'}{n}, \frac{j'}{n}\right) \right| &\leq \epsilon, \\ \left| f_2\left(\frac{i}{n}, \frac{j}{n}\right) - f_2\left(\frac{i'}{n}, \frac{j'}{n}\right) \right| &\leq \epsilon. \end{aligned} \quad (1)$$

Consider the the 4 sets:

$$H^+ = \left\{ (i,j) \in B_n : f_1\left(\frac{i}{n}, \frac{j}{n}\right) - \frac{i}{n} > \epsilon \right\}, \quad (2)$$

$$H^- = \left\{ (i,j) \in B_n : f_1\left(\frac{i}{n}, \frac{j}{n}\right) - \frac{i}{n} < -\epsilon \right\}, \quad (3)$$

$$V^+ = \left\{ (i,j) \in B_n : f_2\left(\frac{i}{n}, \frac{j}{n}\right) - \frac{j}{n} > \epsilon \right\}, \quad (4)$$

$$V^- = \left\{ (i,j) \in B_n : f_2\left(\frac{i}{n}, \frac{j}{n}\right) - \frac{j}{n} < -\epsilon \right\}. \quad (5)$$

Intuitively the points in H^+ under $\mathbf{f}$ are moved further to the right (with increased x coordinate) by more than ϵ . Points in V^- under $\mathbf{f}$ are moved further down (with decreased y coordinate) by more than ϵ . We claim that these sets cannot cover all the vertices of the Hex board. If it were so, then we will have a winner, say Player I with a winning path, linking the East and West boundary frames. Since points of the East boundary have the highest x coordinate, they cannot be moved further to the right. Thus vertices in H^+ are disjoint with the East boundary and similarly vertices in H^- are disjoint with the West boundary. The path must therefore contain vertices from both H^+ and H^- . However for any $(i,j) \in H^+$, $(i',j') \in H^-$ we have

$$\begin{aligned} f_1\left(\frac{i}{n}, \frac{j}{n}\right) - \frac{i}{n} &> \epsilon, \\ -f_1\left(\frac{i'}{n}, \frac{j'}{n}\right) - \frac{i'}{n} &> \epsilon. \end{aligned}$$

Summing the above two inequalities and using (1) we get

$$\frac{i'}{n} - \frac{i}{n} > 2\epsilon.$$

Thus the points (i,j) and (i',j') cannot be adjacent and this contradicts that they are part of a connected path. We have a contradiction.

Remark 7 The algorithm attempts to build a winning path and advances by entering mate triangles. Since the algorithm will not be able to cover the Hex board, partial bridge building should fail at some point, giving a vertex that is outside the

union of sets H^+ , H^- , V^+ , V^- . Hence we reach an approximate fixed point while building the bridge.

An Application of the Algorithm

Consider the continuous map of the unit square into itself given by:

$$\begin{aligned} f_1(x,y) &= \frac{x + \max(-2 + 2x + 6y - 6xy, 0)}{1 + \max(-2 + 2x + 6y - 6xy, 0) + \max(2x - 6xy, 0)} \\ f_2(x,y) &= \frac{y + \max(2 - 6x - 2y + 6xy, 0)}{1 + \max(2 - 6x - 2y + 6xy, 0) + \max(2y - 6xy, 0)}. \end{aligned}$$

With $\epsilon = .05$, we can start with a grid of $\delta = .1$ (hopefully adequate) and find an approximate fixed point. In fact for a spacing of .1 units we have the following iterations. The iterations according to Hex rule passed through the following points with $|f_1(x,y) - x|$ and $|f_2(x,y) - y|$ given by Table 1. Thus the approximate fixed point is $x^* = .4, y^* = .3$.

Extensive Games and Normal Form Reduction

Any game as it evolves can be represented by a rooted tree Γ where the root vertex corresponds to the initial move. Each vertex of the tree represents

Zero-Sum Two Person Games, Table 1 Table giving the Hex building path

(x, y)	$ f_1 - x $	$ f_2 - y $	L
$(.0, .0)$	0	.6667	1
$(.1, 0)$	.0167	.5833	2
$(.1, .1)$	.01228	.4667	2
$(0, .1)$	.0	.53333	1
$(.1, .2)$	.007	.35	2
$(0, .2)$	0	.4	1
$(.1, .3)$	.002	.233	2
$(0, .3)$	0	.26667	1
$(.1, .4)$	.238	.116	1
$(.2, .4)$	.194	.088	1
$(.2, .3)$	.007	.177	2
$(.3, .4)$	.153	.033	1
$(.3, .3)$	.017	.067	2
$(.4, .4)$	.116	0	1
$(.4, .3)$	.0296	0	*

a particular move of a particular player. The alternatives available in any given move are identified with the edges emanating from the vertex that represents the move. If a vertex is assigned to chance, then the game associates a probability distribution with the the descending edges. The terminal vertices are called plays and they are labeled with the payoff to Player I. In zero-sum games Player II's payoff is simply the negative of the payoff to Player I. The vertices for a player are further partitioned into *information sets*. Information sets must satisfy the following requirements:

- The number of edges descending from any two moves within an information set are same.
- No information set intersects the unique unicursal path from the root to any end vertex of the tree in more than one move.
- Any information set which contains a chance move is a singleton.

We will use the following example to illustrate the extensive form representation:

Example 8 Player I has 3 dice in his pocket. Die 1 is a fake die with all sides numbered one. Die 2 is a fake die with all sides numbered two. Die 3 is a genuine unbiased die. He chooses one of the 3 dice secretly, tosses the die once, and announces the outcome to Player II. Knowing the outcome but not knowing the chosen die, Player II tries to guess the die that was tossed. He pays \$1 to Player I if his guess is wrong. If he guesses correctly, he pays nothing to Player I.

The game is represented by the above tree with the root vertex assigned to Player I. The 3 alternatives at this move are to choose the die with all sides 1 or to choose the die with all sides 2 or to choose the unbiased die. The end vertices of these edges descending from the root vertex are moves for chance. The certain outcomes are 1 and 2 if the die is fake. The outcome is one of the numbers 1, ..., 6 if the die chosen is genuine. The other ends of these edges are moves for Player II. These moves are partitioned into information sets V_1 (corresponding to outcome 1), V_2 (corresponding to outcome 2), and singleton information sets V_3, V_4, V_5, V_6 corresponding to outcomes 3, 4, 5 and 6 respectively. Player II must guess the die based


on the information given. If he is told that the game has reached a move in information set V_1 , it simply means that the outcome of the toss is 1. He has two alternatives for each move of this information set. One corresponds to guessing the die is fake and the other corresponds to guessing it is genuine. The same applies to the information set V_2 . If the outcome is in $V_3, \dots, V_6$, the clear choice is to guess the die as genuine. Thus a *pure strategy* (master plan) for Player II is to choose a 2-tuple with coordinates taking the values F or G . Here there are 4 pure strategies for Player II. They are: $(F_1, F_2), (F_1, G), (G, F_2), (G, G)$. For example, the first coordinate of the strategy indicates what to guess when the outcome is 1 and the second coordinate indicates what to guess for the outcome 2. For all other outcomes II guesses the die is genuine (unbiased). The payoff to Player I when Player I uses a pure strategy i and Player II uses a pure strategy j is simply the expected income to Player I when the two players choose i and j *simultaneously*. This can as well be represented by a matrix $A = (a_{ij})$ whose rows and columns are pure strategies and the corresponding entries are the expected payoffs. The payoff matrix A given by

$$A = \begin{matrix} & \begin{matrix} (F_1, F_2) & (F_1, G) & (G, F_2) & (G, G) \end{matrix} \\ \begin{matrix} F_1 \\ F_2 \\ G \end{matrix} & \begin{pmatrix} 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 \\ \frac{1}{3} & \frac{1}{6} & \frac{1}{6} & 0 \end{pmatrix} \end{matrix}.$$

is called the normal form reduction of the original extensive game.

Saddle Point

The normal form of a zero sum two person game has a saddle point when there is a row r and column c such that the entry a_{rc} is the smallest in row r and the largest in column c . By choosing the pure strategy corresponding to row r Player I guarantees a payoff $a_{rc} = \min_j a_{rj}$. By choosing column c , Player II guarantees a loss no more than $\max_i a_{ic} = a_{rc}$. Thus row r and column c are good pure strategies for the two players. In a payoff



Zero-Sum Two Person Games, Fig. 2 Game tree for a single throw with fake or genuine dice

matrix $A = (a_{ij})$ row r is said to *strictly dominate* row t if $a_{rj} > a_{tj}$ for all j . Player I, the maximizer, will avoid row t when it is dominated. If rows r and t are not identical and if $a_{rj} \geq a_{tj}$, then we say that row r *weakly dominates* row t .

Example 9 Player I chooses either 1 or 2. Knowing player I's choice Player II chooses either 3 or 4. If the total T is odd, Player I wins $\$T$ from Player II. Otherwise Player I pays $\$T$ to Player II.

The pure strategies for Player I are simply $\sigma_1 = \text{choose 1}$, $\sigma_2 = \text{choose 2}$. For Player II there are four pure strategies given by: τ_1 : choose 3 no matter what I chooses. τ_2 : choose 4 no matter what I chooses. τ_3 : choose 3 if I chooses 1 and choose 4 if I chooses 2. τ_4 : choose 4 if I chooses 3 1 and choose 3 if I chooses 2. This results in a normal form with payoff matrix A for Player I given by:

$$A = \begin{matrix} & \begin{matrix} (3,3) & (4,4) & (3,4) & (4,3) \end{matrix} \\ \begin{matrix} 1 \\ 2 \end{matrix} & \begin{bmatrix} -4 & 5 & -4 & 5 \\ 5 & -6 & -6 & 5 \end{bmatrix} \end{matrix}$$

Here we can delete column 4 which dominates column 3. We don't have row domination yet. We can delete column 2 as it weakly dominates column 3. Still we have no row domination after these deletions. We can delete column 1 as it weakly dominates column 3. Now we have strict row domination of row 2 by row 1 and we are left with the row 1, column 3 entry = -4 . This is a

saddle point for this game. In fact we have the following:

Theorem 10 The normal form of any zero sum two person game with perfect information admits a saddle point. A saddle point can be arrived at by a sequence of row or column deletions. A row that is weakly dominated by another row can be deleted. A column that weakly dominates another column can be deleted. In each iteration we can always find a weakly or strictly dominated row or a weakly or strictly dominating column to be deleted from the current submatrix.

Mixed Strategy and Minimax Theorem

Zero sum two person games do not always have saddle points in pure strategies. For example, in the game of guessing the die (Example 8) the normal form has no saddle point. Therefore it makes sense for players to choose the pure strategies via a random mechanism. Any probability distribution on the set of all pure strategies for a player is called a *mixed strategy*. In Example 8 a mixed strategy for Player I is a 3-tuple $x = (x_1, x_2, x_3)$ and a mixed strategy for Player II is a 4-tuple $y = (y_1, y_2, y_3, y_4)$. Here x_i is the probability that player I chooses pure strategy i and y_j is the probability that player II chooses pure strategy j .

Since the players play independently and make their choices simultaneously, the expected payoff

to Player I from Player II is $K(x, y) = \sum_i \sum_j a_{ij} x_i y_j$ where a_{ij} are elements in the payoff matrix A .

We call $K(x, y)$ the mixed payoff where players choose mixed strategies x and y instead of pure strategies i and j . Suppose $x^* = (\frac{1}{8}, \frac{1}{8}, \frac{3}{4})$ and $y^* = (\frac{3}{4}, 0, 0, \frac{1}{4})$. Here x^* guarantees Player I an expected payoff of $\frac{1}{4}$ against any pure strategy j of Player II. By the affine linearity of $K(x^*, y)$ in y it follows that Player I has a guaranteed expectation $= \frac{1}{4}$ against any mixed strategy choice of II. A similar argument shows that Player II can choose the mixed strategy $(\frac{3}{4}, 0, 0, \frac{1}{4})$ which limits his maximum expected loss to $\frac{1}{4}$ against any mixed strategy choice of Player I. Thus

$$\max_x K(x, y^*) = \max_y K(x^*, y) = \frac{1}{4}.$$

By replacing the rows and columns with mixed strategy payoffs we have a saddle point in mixed strategies.

Historical Remarks

The existence of a saddle point in mixed strategies for Example 8 is no accident. All finite games have a saddle point in mixed strategies. This important theorem, called the *minimax theorem*, is the very starting point of game theory. While Borel (see under Ville (1938)) considered the notions of pure and mixed strategies for zero sum two person games that have symmetric roles for the players, he was able to prove the theorem only for some special cases. It was von Neumann (1928) who first proved the minimax theorem using some intricate fixed point arguments. While several proofs are available for the same theorem (Gale et al. 1951; Kakutani 1941; Loomis 1946; Nash 1950; Owen 1985); Ville 1938; Weyl 1950, the proofs by Ville and Weyl are notable from an algorithmic point of view. The proofs by Nash and Kakutani allow immediate extension to Nash equilibrium strategies in many person non zero sum games. For zero-sum two person games, optimal strategies and Nash equilibrium strategies coincide. The following is the seminal minimax theorem for matrix games.

Theorem 11 (von Neumann) Let $A = (a_{ij})$ be any $m \times n$ real matrix. Then there exists a pair of

probability vectors $x = (x_1, x_2, \dots, x_m)$ and $y = (y_1, y_2, \dots, y_n)$ such that for a unique constant v

$$\sum_{i=1}^m a_{ij} x_i \geq v \quad j = 1, 2, \dots, n,$$

$$\sum_{j=1}^n a_{ij} y_j \leq v \quad i = 1, 2, \dots, m.$$

The probability vectors x, y are called optimal mixed strategies for the players and the constant v is called the value of the game.

H. Weyl (1950) gave a complete algebraic proof and proved that the value and some pair of optimal strategies for the two players have all of their coordinates lie in the same ordered subfield as the smallest ordered field containing the payoff entries. Unfortunately his proof was non-constructive. It turns out that the minimax theorem can be proved via linear programming in a constructive way which leads to an efficient computational algorithm *a la* the simplex method (Dantzig 1951). The key idea is to convert the problem to dual linear programming problems.

Solving for Value and Optimal Strategies via Linear Programming

Without loss of generality we can assume that the payoff matrix $A = (a_{ij})_{m \times n} > 0$, that is $a_{ij} > 0$ for all (i, j) . Thus we are looking for some v such that:

$$v = \min v_1 \tag{6}$$

such that

$$\sum_{j=1}^n a_{ij} y_j \leq v_1, \tag{7}$$

$$y_1, \dots, y_n \geq 0, \tag{8}$$

$$\sum_{j=1}^n y_j = 1. \tag{9}$$

Since the payoff matrix is positive, any v_1 satisfying the constraints above will be positive, so the problem can be reformulated as

$$\max \frac{1}{v_1} = \max \sum_{j=1}^n \eta_j \quad (10)$$

such that

$$\sum_{j=1}^n a_{ij} \eta_j \leq 1 \quad \text{for all } i, \quad (11)$$

$$\eta_j \geq 0 \quad \text{for all } j. \quad (12)$$

With $A > 0$, the η_j 's are bounded. The maximum of the linear function $\sum_j \eta_j$ is attained at some extreme point of the convex set of constraints (11) and (12). By introducing nonnegative slack variables $s_1, s_2, \dots, s_m$ we can replace the inequalities (11) by equalities (13). The problem reduces to

$$\max \sum_{j=1}^n \eta_j \quad (13)$$

subject to

$$\sum_{j=1}^n a_{ij} \eta_j + s_i = 1, \quad i = 1, 2, \dots, m, \quad (14)$$

$$y_j \geq 0, \quad j = 1, 2, \dots, n, \quad (15)$$

$$s_i \geq 0, \quad i = 1, 2, \dots, m. \quad (16)$$

Of the various algorithms to solve a linear programming problem, the simplex algorithm is among the most efficient. It was first investigated by Fourier (1890). But no other work was done for more than a century. The need for its industrial application motivated active research and lead to the pioneering contributions of Kantorowich (1939) (see a translation in Management Science (Kantorowich 1960)) and Dantzig (1951). It was Dantzig who brought out the earlier investigations of Fourier to the forefront of modern applied mathematics.

Simplex Algorithm

Consider our linear programming problem above. Any solution $\eta = (y_1, \dots, y_n)$, $s = (s_1, \dots, s_m)$ to the above system of equations is called a feasible solution. We could also rewrite the system as

$$\begin{aligned} \eta_1 C^1 + \eta_2 C^2 + \dots + \eta_n C^n + s_1 e^1 + s_2 e^2 + s_m e^m &= \mathbf{1} \\ \eta_1, \eta_2, \dots, \eta_n, s_1, s_2, \dots, s_m &\geq 0. \end{aligned}$$

Here $C^j, j = 1, \dots, n$ are the columns of the matrix A and e^i are the columns of the $m \times m$ identity matrix. The vector $\mathbf{1}$ is the vector with all coordinates unity. With any extreme point $(\eta, s) = (\eta_1, \eta_2, \dots, \eta_n, s_1, \dots, s_m)$ of the convex polyhedron of feasible solutions one can associate with it a set of m linearly independent columns, which form a basis for the column span of the matrix (A, I) . Here the coefficients η_j and s_i are equal to zero for coordinates *other* than for the specific m linearly independent columns. By slightly perturbing the entries we can assume that any extreme point of feasible solutions has *exactly* m positive coordinates. Two extreme feasible solutions are called *adjacent* if the associated bases differ in exactly one column.

The key idea behind the simplex algorithm is that an extreme point $P = (\eta^*, s^*)$ is an optimal solution if and only if there is no adjacent extreme point Q for which the objective function has a higher value. Thus when the algorithm is initiated at an extreme point which is not optimal, there must be an adjacent extreme point that strictly improves the objective function. In each iteration, a column from outside the basis replaces a column in the current basis corresponding to an adjacent extreme point. Since there are $m + n$ columns in all for the matrix (A, I) , and in each iteration we have strict improvement by our non-degeneracy assumption on the extreme points, the procedure must terminate in a finite number of steps resulting in an optimal solution.

Example 12 Players I and II simultaneously show either 1 or 2 fingers. If T is the total number of fingers shown then Player I receives from Player II $\$T$ when T is odd and loses $\$T$ to Player II when T is even.

The payoff matrix is given by

$$A = \begin{bmatrix} -2 & 3 \\ 3 & -4 \end{bmatrix}.$$

Add 5 to each entry to get a new payoff matrix with all entries strictly positive. The new game is strategically same as A .

$$\begin{bmatrix} 3 & 8 \\ 8 & 1 \end{bmatrix}.$$

The linear programming problem is given by
 $\max 1 \cdot y_1 + 1 \cdot y_2 + 0 \cdot s_1 + 0 \cdot s_2$ such that

$$\begin{bmatrix} 3 & 8 & 1 & 0 \\ 8 & 1 & 0 & 1 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ s_1 \\ s_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

We can start with the trivial solution $(0, 0, 1, 1)^T$. This corresponds to the basis e^1, e^2 with $s_1 = s_2 = 1$. The value of the objective function is 0. If we make $y_2 > 0$, then the value of the objective function can be increased. Thus we look for a solution to

$$s_1 = 0, \quad s_2 > 0, \quad y_2 > 0$$

satisfying the constraints

$$\begin{aligned} 8y_2 + 0s_2 &= 1 \\ y_2 + s_2 &= 1 \end{aligned}$$

or to

$$s_2 = 0, \quad s_1 > 0, \quad y_2 > 0$$

satisfying the constraints

$$\begin{aligned} 8y_2 + s_1 &= 1 \\ y_2 + 0s_1 &= 1 \end{aligned}$$

Notice that $y_2 = \frac{1}{8}$, $s_2 = \frac{7}{8}$ is a solution to the first system and that the second system has no nonnegative solution. The value of the objective function at this extreme solution is $\frac{1}{8}$. Now we look for an adjacent extreme point. We find that $y_1 = \frac{7}{61}$, $y_2 = \frac{5}{61}$ is such a solution. The procedure terminates because no adjacent solution with $y_1 > 0$, $s_1 > 0$ or $y_2 > 0$, $s_1 > 0$ or $y_1 > 0$, $s_2 > 0$ or $y_2 > 0$, $s_2 > 0$ if any has higher objective function value. The algorithm terminates with the optimal value of $\frac{1}{v_1} = \frac{12}{61}$. Thus the value of the modified game is $\frac{61}{12}$, and the value of the original game is $\frac{61}{12} - 5 = \frac{1}{12}$. A good strategy

for Player II is obtained by normalizing the optimal solution of the linear program, it is $\eta_1 = \frac{7}{12}$, $\eta_2 = \frac{5}{12}$. Similarly, from the dual linear program we can see that the strategy $\xi_1 = \frac{7}{12}$, $\xi_2 = \frac{5}{12}$ is optimal for Player I.

Fictitious Play

Though optimal strategies are not easily found, even naive players can learn to steer their average payoff towards the value of the game from past plays by certain iterative procedures. This learning procedure is known as *fictitious play*. The two players make their next choice under the assumption that the opponent will continue to choose pure strategies at the same frequencies as what he/she did in the past. If $x^{(n)}, y^{(n)}$ are the empirical mixed strategies used by the two players in the first n rounds, then in round $n+1$ Player I pretends that Player II will continue to use $y^{(n)}$ in the future and selects any row i^* such that

$$\sum_j a_{i^*j} y_j^{(n)} = \max_i \sum_j a_{ij} y_j^{(n)}.$$

The new empirical mixed strategy is given by

$$x^{(n+1)} = \frac{1}{n+1} I_{i^*} + \frac{n}{n+1} x^{(n)}.$$

(Here I_{i^*} is the degenerate choice of pure strategy i^* .) This intuitive learning procedure was proposed by Brown (1951) and the following convergence theorem was proved by Robinson (1951).

Theorem 13
 $\lim_n \min_j \sum_i a_{ij} x_i^{(n)} = \lim_n \min_i \sum_j a_{ij} y_j^{(n)} = v.$

We will apply the fictitious play algorithm to the following example and get a bound on the value.

Example 14 Player I picks secretly a card of his choice from a deck of three cards numbered 1, 2, and 3. Player II proceeds to guess player I's choice. After each guess player I announces

player II's guess as "High", "Low" or "Correct" as the case may be. The game continues till player II guesses player I's choice correctly. Player II pays to Player I \$ N where N is the number of guesses he made.

The payoff matrix is given by

$$A = \begin{matrix} & \begin{matrix} (1,2) & (1,3) & (2) & (3,1) & (3,2) \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{pmatrix} 1 & 1 & 2 & 2 & 3 \\ 2 & 3 & 1 & 3 & 2 \\ 3 & 3 & 2 & 2 & 1 \end{pmatrix} \end{matrix}.$$

Here the row labels are possible cards chosen by Player I, and the column labels are pure strategies for Player II. For example, the pure strategy (1, 3) for Player II, means that 1 is the first guess and if 1 is incorrect then 3 is the second guess. The elements of the matrix are payoffs to Player I. We can use cumulative total, instead average for the players to make their next choice of row or column based on the totals. We choose the row or column with the least index in case more than one row or one column meets the criterion. The total for the first 10 rounds using fictitious play is given in Table 2.

The bold entries give the approximate lower and upper bounds for the total payoff in 10 rounds giving $1.6 \leq v \leq 1.9$.

Remark 15 Fictitious play is known to have a very poor rate of convergence to the value. While it works for all zero sum two person games, it fails to extend to Nash equilibrium payoffs in bimatrix games even when the game has a

unique Nash equilibrium. It extends only to some very special classes like 2×2 bimatrix games and to the so called potential games.

(See Miyasawa (1961), Shapley (1964), Monderer and Shapley (1996), Krishna and Sjoestrom (1998), Berger (2007)).

Search Games

Search games are often motivated by military applications. An object is hidden in space. While the space where the object is hidden is known, the exact location is unknown. The search strategy consists of either targeting a single point of the space and paying a penalty when the search fails or continue the search till the searcher gets closer to the hidden location. If the search consists of many attempts, then a pure strategy is simply a function of the search history so far. Example 14 is a typical search game. The following are examples of some simple search games that have unexpected turns with respect to the value and optimal strategies.

Example 16 A pet shop cobra of known length $t < 1$ escapes out of the shop and has settled in a nearby tree somewhere along a particular linear branch of unit length. Due to the camouflage, the exact location $[x, x + t]$ where it has settled on the branch is unknown. The shop keeper chooses a point y of his choice and aims a bullet at the point y . In spite of his 100% accuracy the cobra will escape permanently if his targeted point y is outside the settled location of the cobra.

Zero-Sum Two Person Games, Table 2

Row choices	Total so far					Column choices	Total so far		
R_1	1	1	2	2	3	C_1	1	2	3
R_3	4	3	4	3	4	C_2	2	5	5
R_2	6	6	5	6	6	C_3	4	6	7
R_3	9	8	7	7	7	C_3	6	7	9
R_3	12	10	9	8	8	C_4	8	10	10
R_2	15	12	11	9	9	C_4	10	13	11
R_2	17	15	12	12	11	C_5	13	15	12
R_2	19	18	13	15	13	C_3	15	16	14
R_2	21	21	14	18	15	C_3	17	17	16
R_1	22	22	16	20	18	C_3	19	18	18

We treat this as a game between the cobra (Player I) and the shop keeper (Player II). Let the probability of survival be the payoff to the cobra. Thus

$$K(x, y) = \begin{cases} 1 & \text{if } y < x, \text{ or } y > x + t \\ 0 & \text{otherwise.} \end{cases}$$

The pure strategy spaces are $0 \leq x \leq 1 - t$ for the snake and $0 \leq y \leq 1$ for the shop keeper. It can be shown that the game has no saddle point and has optimal mixed strategies. The value function $v(t)$ is a discontinuous function of t . In case $\frac{1}{t}$ is an integer n then, a good strategy for the snake is to hide along $[0, t]$, or $[t, 2t]$ or $[(n - 1)t, 1]$ chosen with equal chance. In this case the optimal strategy for the shop keeper is to choose a random point in $[0, 1]$. The value is $1 - \frac{1}{n}$. In case $\frac{1}{t}$ is a fraction, let $n = \lceil \frac{1}{t} \rceil$ then the optimal strategy for the snake is to hide along $[0, t]$, or $[t, 2t]$, ... or $[(n - 1)t, nt]$. An optimal strategy for the shop keeper is to shoot at one of the points $\frac{1}{n+1}, \frac{2}{n+1}, \dots, \frac{n}{n+1}$ chosen at random.

Example 17 While mowing the lawn a lady suddenly realizes that she has lost her diamond engagement ring some where in her lawn. She has maximum speed s and will be able to locate the diamond ring from its glitter if she is sufficiently close to, say within a distance ϵ from the ring. What is an optimal search strategy that minimizes her search time.

If we treat Nature as a player against her, she is playing a zero sum two person game where Nature would find pleasure in her delayed success in finding the ring.

Search Games on Trees

The following is an elegant search game on a tree. For many other search games the readers can refer to the monographs by Gal (1980) and Alpern and Gal (2003). Also see (Reijnders and Potters 1993).


Example 18 A bird has to look for a suitable location to build its nest for hatching eggs and protecting them against predator snakes. Having identified a large tree with a single predator snake in the neighborhood, the bird has to further decide where to build its nest on the chosen tree. The chance for

the survival of the eggs is directly proportional to the distance the snake travels to locate the nest.

While birds and snakes work out their strategies based on instinct and evolutionary behavior, we can surely approximate the problem by the following zero sum two person search game. Let $T = (X, E)$ be a finite tree with vertex set X and edge set E . Let $O \in X$ be the root vertex. A hider hides an object at a vertex x of the tree. A searcher starts at the root and travels along the edges of the tree such that the path traced covers all the terminal vertices. The search ends as soon as the searcher crosses the hidden location and the payoff to the hider is the distance traveled so far.

By a simple domination argument we can as well assume that the optimal hiding locations are simply the terminal vertices.

Theorem 19 The search game has value and optimal strategies. The value coincides with the sum of all edge lengths. Any optimal strategy for the hider will necessarily restrict to hide at one of the terminal vertices. Let the least distance traveled to exhaust all end vertices one by one correspond to a permutation σ of the end vertices in the order $w_1, w_2, \dots, w_k$. Let σ^{-1} be its reverse permutation. Then an optimal strategy for the searcher is to choose one of these two permutations by the toss of a coin. The hider has a unique optimal mixed strategy that chooses each end vertex with positive probability.



Zero-Sum Two Person Games, Fig. 3 Bird trying to hide at a leaf and snake chasing to reach the appropriate leaf via optimal Chinese postman route starting at root O and ending at O

Suppose the tree is a path with root O and a single terminal vertex x . Since the search begins at O , the longest trip is possible only when hider hides at x and the theorem holds trivially. In case the tree has just two terminal vertices besides the root vertex, the possible hiding locations are say, O, x_1, x_2 with edge lengths a_1, a_2 . The possible searches are via paths: $O \rightarrow x_1 \rightarrow O \rightarrow x_2$ abbreviated Ox_1Ox_2 or $O \rightarrow x_2 \rightarrow O \rightarrow x_1$, abbreviated Ox_2Ox_1 . The payoff matrix can be written as

$$\begin{array}{cc} Ox_1Ox_2 & Ox_2Ox_1 \\ \begin{array}{c} x_1 \\ x_2 \end{array} \begin{bmatrix} a_1 & 2a_2 + a_1 \\ 2a_1 + a_2 & a_2 \end{bmatrix} \end{array}.$$

The value of this game is $a_1 + a_2 =$ sum of the edge lengths. We can use an induction on the number of subtrees to establish the value as the sum of edge lengths. We will use an example to just provide the intuition behind the formal proof.

Given any permutation τ of the end vertices (leaves), of the above tree let P be the shortest path from the root vertex that travels along the leaves in that order and returns to the root. Let the reverse path be P^{-1} . Observe that it will cover all edges twice. Thus if the two paths P and P^{-1} are chosen with equal chance by the snake, the average distance traveled by the snake when it locates the bird's nest at an end vertex will be independent of the particular end vertex. For example along the closed path $O \rightarrow t \rightarrow d \rightarrow t \rightarrow a \rightarrow t \rightarrow x \rightarrow b \rightarrow x \rightarrow c \rightarrow x \rightarrow t$ the distance traveled by the snake to reach leaf e is $(3 + 7 + 7 + \dots + 8 + 3) = 74$. If the snake travels along the reverse path to reach e the distance traveled is $(9 + 5 + 5 + \dots + 5 + 4 + 3) = 66$. For example if it is to reach the vertex d then via path P it is $(3 + 7)$. Via P^{-1} it is to make travel to e and travel from e to d by the reverse path. This is $(66 + 3 + \dots + 6 + 6 + 7) = 130$. Thus in both cases the average distance traveled is 70.

The average distance is the same for every other leaf when P and P^{-1} are used. The optimal Chinese postman route can allow all permutations subject to permuting any leaf of any subtree only among themselves. Thus the subtree rooted at t has leaves a, b, c, d and the subtree rooted at u has leaves e, f, g, h, j . For example while permuting a, b, c, d only among themselves we have the further restriction that between b, c we cannot allow insertion of a or d . For example a, b, c, d and a, d, c, b are acceptable permutations, but not a, b, d, c . It can never be the optimal permuting choice. The same way it applies to the tree rooted at u . For example h, j, e, g, f is part of the optimal Chinese postman route, but h, g, j, e, f is not. We can think of the snake and bird playing the game as follows: The bird chooses to hide in a leaf of the subgame G_t rooted at t or at a leaf of the subgame G_u rooted at u . These leaves exhaust all leaves of the original game. The snake can restrict to only the optimal route of each subgame. This can be thought of as a 2×2 game where the strategies for the two players (bird) and snake are:

Bird:

Strategy 1: Hide optimally in a leaf of G_t

Strategy 2: Hide optimally in a leaf of G_u .

Snake:

Strategy 1: Search first the leaves of G_t along the optimal Chinese postman route of G_t and then search along the leaves of G_u .

Strategy 2: Search first the leaves of G_u along the optimal Chinese postman route and then search the leaves of G_t along the optimal postman route. The expected outcome can be written as the following 2×2 game. (Here $v(G_t)$, $v(G_u)$ are the values of the subgames rooted at t, u respectively.)

$$\begin{array}{cc} G_t G_u G_t G_u \\ \begin{array}{c} G_t \\ G_u \end{array} \begin{bmatrix} [3 + v(G_t)] & 2[9 + v(G_u) + [3 + v(G_t)]] \\ 2[3 + v(G_t)] + [9 + v(G_u)] & [9 + v(G_u)] \end{bmatrix} \end{array}$$

Observe that the 2×2 game has no saddle point and hence has value $3 + 9 + v(G_t) + v(G_u)$. By induction we can assume $v(G_t) = 24$, $v(G_u) = 34$. Thus the value of this game is 70. This is also the sum of the edge lengths of the game tree. An optimal strategy for the bird can be recursively determined as follows.

Umbrella Folding Algorithm

Ladies, when storing umbrellas inside their handbag shrink the central stem of the umbrella and then the stems around all in one stroke. We can mimic a somewhat similar procedure also for our above game tree. We simultaneously shrink the edges $[xc]$ and $[xb]$ to x . In the next round $\{a, x, d\}$ edges $[a, t]$, $[x, t]$, $[d, t]$ can be simultaneously shrunk to t and so on till the entire tree is shrunk to the root vertex O . We do know that the optimal strategy for the bird when the tree is simply the subtree with root x and with leaves b, c is given by $p(b) = \frac{4}{(4+2)}$, $p(c) = \frac{2}{(4+2)}$. Now for the subtree with vertex t and leaves $\{a, b, c, d\}$, we can treat this as collapsing the previous subtree to x and treat stem length of the new subtree with vertices $\{t, a, x, d\}$ as though the three stems $[ta]$, $[tx]$, $[td]$ have lengths 6, $5 + (4 + 2)$, 7. We can check that for this subtree game the leaves a, x, d are chosen with probabilities $p(a) = \frac{6}{(6+9+7)}$, $p(x) = \frac{9}{(6+9+7)}$, $p(d) = \frac{7}{(6+9+7)}$. Thus the optimal mixed strategy for the bird for choosing leaf b for our original tree game is to pass through vertices t, x, b and is given by the product $p(t)p(x)p(b)$. We can inductively calculate these probabilities.

Completely Mixed Games and Perron's Theorem on Positive Matrices

A mixed strategy x for player I is called *completely mixed* if it is strictly positive ($x > 0$). A matrix game A is *completely mixed* if and only all optimal mixed strategies for Player I and Player II are completely mixed. The following elegant theorem was proved by Kaplanski (1945).

Theorem 20 A matrix game A with value v is completely mixed if and only if

1. The matrix is square.
2. The optimal strategies are unique for the two players.
3. If $v \neq 0$, then the matrix is nonsingular.
4. If $v = 0$, then the matrix has rank $n - 1$ where n is the order of the matrix.

The theory of completely mixed games is a useful tool in linear algebra and numerical analysis (Bapat and Raghavan 1997). The following is a sample application of this theorem.

Theorem 21 (Perron 1909) Let A be any $n \times n$ matrix with positive entries. Then A has a positive eigenvalue with a positive eigenvector which is also a simple root of the characteristic equation.

Proof Let I be the identity matrix. For any $\lambda > 0$, the maximizing player prefers to play the game A rather than the game $A - \lambda I$. The payoff gets worse when the diagonal entries are reached. The value function $v(\lambda)$ of $A - \lambda I$ is a non-increasing continuous function. Since $v(0) > 0$ and $v(\lambda) < 0$ for large λ we have for some $\lambda_0 > 0$ the value of $A - \lambda_0 I$ is 0. Let y be optimal for player II, then $(A - \lambda_0 I)y \leq 0$ implies $0 < Ay \leq \lambda_0 y$. That is 0. Since the optimal y is completely mixed, for any optimal x of player I, we have $(A - \lambda_0 I)x = 0$. Thus $x > 0$ and the game is completely mixed. By (2) and (4) if $(A - \lambda_0 I)u = 0$ then u is a scalar multiple of y and so the eigenvector y is geometrically simple. If $B = A - \lambda_0 I$, then B is singular and of rank $n - 1$. If (B_{ij}) is the cofactor matrix of the singular matrix B then $\sum_j b_{ij} B_{kj} = 0$, $i = 1, \dots, n$. Thus row k of the cofactor matrix is a scalar multiple of y . Similarly each column of B is a scalar multiple of x . Thus all cofactors are of the same sign and are different from 0. That is

$$\frac{d}{d\lambda} \det(A - \lambda I) |_{\lambda_0} = \sum_i B_{ii} \neq 0.$$

Thus λ_0 is also algebraically simple. See (Bapat and Raghavan 1997) for the most general

extensions of this theorem to the theorems of Perron and Frobenius and to the theory of M -matrices and power positive and polynomially matrices).

Behavior Strategies in Games with Perfect Recall

Consider any extensive game Γ where the unique unicursal path from an end vertex w to the root x_0 intersects two moves x and y of say, Player I. We say $x < y$ if the the unique path from y to x_0 is via move x . Let $U \ni x$ and $V \ni y$ be the respective information sets. If the game has reached a move $y \in V$, Player I will know that it is his turn and the game has progressed to *some* move in V . The game is said to have *perfect recall* if each player can remember all his past moves and the choices made in those moves. For example if the game has progressed to a move of Player I in the information set V he will remember the specific alternative chosen in any earlier move. A move x is *possible* for Player I with his pure strategy π_1 , if for some suitable pure strategy π_2 of Player II, the move x can be reached with positive probability using π_1, π_2 . An information set U is *relevant* for a pure strategy π_1 , for Player I, if some move $x \in U$ is possible with π_1 . Let Π_1, Π_2 be pure strategy spaces for players I and II.

Let $\mu_1 = \{q_{\pi_1, \pi_1 \in \Pi_1}\}$ be any mixed strategy for Player I. The information set U for Player I is *relevant for the mixed strategy* μ_1 if for some $q_{\pi_1} > 0$, U is relevant for π_1 . We say that the information set U for Player I is not relevant for the mixed strategy μ_1 if for all $q_{\pi_1} > 0$, U is not relevant for π_1 . Let

$$S_v = \{\pi_1 : U \text{ is relevant for } \pi_1 \text{ and } \pi_1(U) = v\},$$

$$S = \{\pi_1 : U \text{ is relevant for } \pi_1\},$$

$$T = \{\pi_1 : U \text{ is not relevant for } \pi_1 \text{ and } \pi_1(U) = v\}.$$

The behavior strategy induced by a mixed strategy pair (μ_1, μ_2) at an information set U for Player I is simply the conditional probability of choosing alternative v in the information set U ,

given that the game has progressed to a move in U , namely

$$\beta_1(U, v) = \begin{cases} \frac{\sum_{\pi_1 \in S_v} q_{\pi_1}}{\sum_{\pi_1 \in S} q_{\pi_1}} & \text{if } U \text{ is relevant for } \mu_1, \\ \sum_{\pi_1 \in T} q_{\pi_1} & \text{if } U \text{ is not relevant for } \mu_1. \end{cases}$$

The following theorem of Kuhn (1953) is a consequence of the assumption of perfect recall.

Theorem 22 Let μ_1, μ_2 be mixed strategies for players I and II respectively in a zero sum two person finite game Γ of perfect recall. Let β_1, β_2 be the induced behavior strategies for the two players. Then the probability of reaching any end vertex w using μ_1, μ_2 coincides with the probability of reaching w using the induced behavior strategy β_1, β_2 . Thus in zero-sum two person games with perfect recall, players can play optimally by restricting their strategy choices just to behavior strategies.

The following analogy may help us understand the advantages of behavior strategies over mixed strategies. A book has 10 pages with 3 lines per page. Someone wants to glance through the book reading just 1 line from each page. A master plan (pure strategy) for scanning the book consists of choosing one line number for each page. Since each page has 3 lines, the number of possible plans is 3^{10} . Thus the set of mixed strategies is a set of dimension $3^{10} - 1$. There is another randomized approach for scanning the book. When page i is about to be scanned choose line 1 with probability x_{i1} , line 2 with probability x_{i2} and line 3 with probability x_{i3} . Since for each i we have $x_{i1} + x_{i2} + x_{i3} = 1$ the dimension of such a strategy space is just 20. Behavior strategies are easier to work with. Further Kuhn's theorem guarantees that we can restrict to behavior strategies in games with perfect recall.

In general if there are k alternatives at each information set for a player and if there are n information sets for the player, the dimension of the mixed strategy space is $k^n - 1$. On the other hand the dimension of the behavior strategy space is simply $n(k - 1)$. Thus while the dimension of

mixed strategy space grows exponentially the dimension of behavior strategy space grows linearly. The following example will illustrate the advantages of using behavior strategies.

Example 23 Player I has 7 dice. All but one are fake. Fake die F_i has the same number i on all faces $i = 1, \dots, 6$. Die G is the ordinary unbiased die. Player I selects one of them secretly and announces the outcome of a single toss of the die to player II. It is Player II's turn to guess which die was selected for the toss. He gets no reward for correct guess but pays \$1 to Player I for any wrong guess.

Player I has 7 pure strategies while Player II has 2^6 pure strategies. As an example the pure strategy (F_1, G, G, F_4, G, F_6) for Player II is one which guesses the die as fake when the outcome revealed is 1 or 4 or 6, and guesses the die as genuine when the outcome is 2 or 3 or 5. The normal form game is a payoff matrix of size 7×64 . For example if G is chosen by Player I, and (F_1, G, G, F_4, G, F_6) is chosen by Player II, the expected payoff to Player I is $\frac{1}{6}[1 + 0 + 0 + 1 + 0 + 1] = \frac{1}{2}$. If F_2 is chosen by Player I, the expected payoff is 1 against the above pure strategy of Player II. Now Player II can use the following behavior strategy. If the outcome is i , then with probability q_i he can guess that the die is genuine and with probability $(1 - q_i)$ he can guess that it is from the fake die F_i . The expected behavioral payoff to Player I when he chooses the genuine die with probability p_0 and chooses the fake die F_i with probability p_i , $i = 1, \dots, 6$ is given by

$$K(p, q) = p_0 \frac{1}{6} \sum_{i=1}^6 (1 - q_i) + \sum_{i=1}^6 p_i q_i.$$

Collecting the coefficients of q_i 's, we get

$$K(p, q) = \sum_{i=1}^6 q_i \left[p_i - \frac{1}{6} p_0 \right] + p_0.$$

By choosing $p_i - \frac{1}{6} p_0 = 0$, we get $p_1 = p_2 = \dots, p_6 = \frac{1}{6} p_0$. Thus $p = (\frac{1}{2}, \frac{1}{12}, \frac{1}{12}, \frac{1}{12}, \frac{1}{12}, \frac{1}{12}, \frac{1}{12})$. For this

mixed strategy for Player I, the payoff to Player I is independent of Player II's actions. Similarly, we can rewrite $K(p, q)$ as a function of p_i 's for $i = 1, \dots, 6$ where

$$K(p, q) = \sum_{k \neq 0} p_k \left[q_k - \frac{1}{6} \sum_{r=1}^6 (1 - q_r) \right] + \frac{1}{6} \left(\sum_{k=1}^6 (1 - q_k) \right).$$

This expression can be made independent of p_i 's by choosing $q_i = \frac{1}{2}$, $i = 1, \dots, 6$. Since the behavioral payoff for these behavioral strategies is $\frac{1}{2}$, the value of the game is $\frac{1}{2}$, which means that Player I cannot do any better than $\frac{1}{2}$ while Player II is able to limit his losses to $\frac{1}{2}$.

Efficient Computation of Behavior Strategies

Introduction

In our above example with one genuine and six fake dice, we used Kuhn's theorem to narrow our search among optimal behavior strategies. Our success depended on exploiting the inherent symmetries in the problem. We were also lucky in our search when we were looking for optimals among equalizers.

From an algorithmic point of view, this is not possible with any arbitrary extensive game with perfect recall. While normal form is appropriate for finding optimal mixed strategies, its complexity grows exponential with the size of the vertex set. The payoff matrix in normal form is in general not a sparse matrix (a sparse matrix is one which has very few nonzero entries) a key issue for data storage and computational accuracies. By sticking to the normal form of a game with perfect recall we cannot take full advantage of Kuhn's theorem in its drastically narrowed down search for optimals among behavior strategies. A more appropriate form for these games is the *sequence form* (von Stengel 1996) and *realization probabilities* to be described below. The behavioral strategies that induce the *realization probabilities* grow only

linearly in the size of the terminal vertex set. Another major advantage is that the sequence form induces a sparse matrix. It has at most as many non-zero entries as the number of terminal vertices or plays.

Sequence Form

When the game moves to an information set U_1 of say, player I, the perfect recall condition implies that wherever the true move lies in U_1 , the player knows the actual alternative chosen in any of the past moves. Let σ_{u_1} denote the sequence of alternatives chosen by Player I in his *past moves*. If no past moves of player I occurs we take $\sigma_{u_1} = \emptyset$. Suppose in U_1 player I selects an action “ c ” with behavioral probability $\beta_1(c)$ and if the outcome is c the new sequence is $\sigma_{u_1} \cup c$. Thus any *sequence* s_1 for player I is the string of choices in his moves along the partial path from the initial vertex to any other vertex of the tree. Let S_0, S_1, S_2 be the set of all sequences for Nature (via chance moves), Player I and Player II respectively. Given behavior strategies $\beta_0, \beta_1, \beta_2$ Let

$$r_i(s_i) = \prod_{c \in s_i} \beta_i(c), \quad i = 0, 1, 2.$$

The functions: $r_i : S_i \rightarrow R: i = 0, 1, 2$ satisfy the following conditions

$$r_i(\emptyset) = 1 \quad (17)$$

$$r_i(\sigma_{u_i}) = \sum_{c \in A(U_i)} r_i(\sigma_{u_i}, c), \quad i = 0, 1, 2 \quad (18)$$

$$r_i(s_i) \geq 0 \text{ for all } s_i, i = 0, 1, 2. \quad (19)$$

Conversely given any such realization functions r_1, r_2 we can define behavior strategies, β_1 say for player I, by

$$\beta_1(U_1, c) = \frac{r_1(\sigma_{u_1} \cup c)}{r_1(\sigma_{u_1})} \text{ for } c \in A(U_1), \\ \text{and } r_1(\sigma_{u_1}) > 0.$$

When $r_1(\sigma_{u_1}) = 0$ we define $\beta_1(U_1, c)$ arbitrarily so that $\sum_{c \in A(U_1)} \beta_1(U_1, c) = 1$. If the ter-

minal payoff to player I at terminal vertex ω is $h(\omega)$, by defining $h(a) = 0$ for all nodes a that are not terminal vertices, we can easily check that the behavioral payoff

$$H(\beta_1, \beta_2) = \sum_{s \in S} h(s) \prod_{i=0}^2 r_i(s_i).$$

When we work with realization functions $r_i, i = 1, 2$ we can associate with these functions the sequence form of payoff matrix whose rows correspond to sequence $s_1 \in S_1$ for Player I and columns correspond to sequence $s_2 \in S_2$ for Player II and with payoff matrix

$$K(s_1, s_2) = \sum_{s_0 \in S_0} h(s_0, s_1, s_2).$$

Unlike the mixed strategies we have more constraints on the sequences r_1, r_2 for each player given by the linear constraints above. It may be convenient to denote the sequence functions r_1, r_2 by vectors x, y respectively. The vector x has $|S_1|$ coordinates and vector y has $|S_2|$ coordinates. The constraints on x and y are linear given by $Ex = e, Fy = f$ where the first row is the unit vector $(1, 0, \dots, 0)$ of appropriate size in both E and F . If u_1 is the collection of information sets for player I then the number of rows in E is $1 + |u_1|$. Similarly the number of rows in F is $1 + |u_2|$. Except for the first row, each row has the starting entry as -1 and some 1's and 0's. Consider the following extensive game with perfect recall.

The set of sequences for player I is given by $S_1 = \{\emptyset, l, r, L, R\}$. The set of sequences for player II is given by $S_2 = \{\emptyset, c, d\}$. The sequence form payoff matrix is given by

$$K(s_1, s_2) = A = \begin{bmatrix} * & * & * \\ 0 & 0 & 0 \\ 0 & 1 & -1 \\ 0 & -2 & 4 \\ 1 & 0 & 0 \end{bmatrix}.$$

The constraint matrices E and F are given by

$$E = \begin{bmatrix} 1 & & & \\ -1 & 1 & 1 & \\ -1 & & & 1 & 1 \end{bmatrix}, \quad F = \begin{bmatrix} 1 & & \\ -1 & 1 & 1 \end{bmatrix}.$$

Since no end vertex corresponds $s_1 = \emptyset$, for Player I, the first row of A is identically 0 and so it is represented by * entries.

We are looking for a pair of vectors x^*, y^* that y^* is the best reply vector y which minimizes $(x^* Ay)$ among all vectors satisfying $Fy = f, y \geq 0$. Similarly the best reply vector is x^* where it maximizes (x, Ay^*) subject to $E^T x = e, x \geq 0$. The duals to these two linear programming problems are

$$\begin{array}{ll} \max(fq) & \min(ep) \\ \text{such that} & \text{such that} \\ F^T q \leq A^T x^* & Ep \geq Ay^* \\ q \text{ unrestricted.} & p \text{ unrestricted.} \end{array}$$

Since $E^T x^* = e$, and $x^* \geq 0$, $Fy^* = f$, and $y^* \geq 0$ these two problems can as well be viewed as the dual linear programs


Primal:

$$\begin{array}{l} \max(fq) \\ \text{such that} \end{array}$$

$$\begin{bmatrix} F^T & -A^T \\ 0 & -E^T \end{bmatrix} \begin{bmatrix} q \\ x \end{bmatrix} \begin{bmatrix} \leq 0 \\ = e \end{bmatrix}, \quad x \geq 0, q \text{ unrestricted.}$$

Dual:

$$\min(ep)$$



Zero-Sum Two Person Games, Fig. 4

such that

$$\begin{bmatrix} F & 0 \\ -A & E \end{bmatrix} \begin{bmatrix} y \\ p \end{bmatrix} \begin{bmatrix} = f \\ \geq 0 \end{bmatrix}, \quad y \geq 0, p \text{ unrestricted.}$$

We can essentially prove that:

Theorem 24 The optimal behavior strategies of a zero sum two person game with perfect recall can be reduced to solving for optimal solutions of dual linear programs induced by its sequence form. The linear program has a size which in its sparse representation is linear in the size of the game tree.

General Minimax Theorems

The minimax theorem of von Neumann can be generalized if we notice the possible limitations for extensions.

Example 25 Players I and II choose secretly positive integers i, j respectively. The payoff matrix is given by

$$a_{ij} = \begin{cases} 1 & \text{if } i > j \\ -1 & \text{if } i < j \end{cases}$$

the value of the game does not exist.

The boundedness of the payoff is essential and the following extension holds.

S-games

Given a closed bounded set $S \subseteq \mathbf{R}^m$, let Players I and II play the following game. Player II secretly selects a point $s = (s_1, \dots, s_m) \in S$. Knowing the set S but not knowing the point chosen by Player II, Player I selects a coordinate $i \in \{1, 2, \dots, m\}$. Player I receives from Player II an amount s_i .

Theorem 26 Given that S is a compact subset of $\mathbf{R}^m$, every S-game has a value and the two players have optimal mixed strategies which use at most m pure strategies. If the set S is also convex, Player II has an optimal pure strategy.

Proof Let T be the convex hull of S . Here T is also compact. Let

$$v = \min_{t \in T} \max_i t_i = \max_i \min_{t \in T} t_i.$$

The compact convex set T and the open convex set $G = \{s : \max_i s_i < v\}$ are disjoint. By the weak separation theorem for convex sets there exists a $\xi \neq 0$ and constant c such that

$$\begin{aligned} \text{for all } s \in G, \quad (\xi, s) &\leq c \text{ and} \\ \text{for all } t \in T, \quad (\xi, t) &\geq c. \end{aligned}$$

Using the property that $\mathbf{v} = (v, v, \dots, v) \in \overline{G}$ and $t^* \in T \cap \overline{G}$, we have $(\xi, t^*) = c$. For any $u \geq 0$, $t^* - u \in \overline{G}$. Thus $(\xi, t^* - u) \leq c$. That is $(\xi, u) \geq 0$. We can assume ξ is a probability vector, in which case

$$(\xi, \mathbf{v}) = v \leq c = (\xi, t^*) \leq \max_i t_i^* = v.$$

Now Player II has $t^* = \sum_j \mu_j x^j$ a mixed strategy which chooses x^j with probability μ_j . Since t^* is a boundary point of $T = \text{con} S$, by the Caratheodary theorem for convex hulls, the convex combination above involves at most m points of S . Hence the theorem. See (Parthasarathy and Raghavan 1971).

Geometric Consequences

Many geometric theorems can be derived by using the minimax theorem for S -games. Here we give as an example the following theorem of Berge (1963) that follows from the theorem on S -games.

Theorem 27 Let S_i , $i = 1, \dots, m$ be compact convex sets in $\mathbf{R}^n$. Let them satisfy the following two conditions.

1. $S = \bigcup_{i=1}^m S_i$ is convex.
 2. $\bigcap_{i \neq j} S_i \neq \emptyset$, $j = 1, 2, \dots, m$.
- Then $\bigcap_{i=1}^m S_i \neq \emptyset$.

Proof Suppose Player I secretly chooses one of the sets S_i and Player II secretly chooses a point $x \in S$. Let the payoff to I be $d(S_i, x)$ where d is the distance of the point x from the set S_i . By our S -game arguments we have for some probability vector ξ , and mixed strategy $\mu = (\mu_1, \dots, \mu_m)$

$$\sum_i \xi_i d(S_i, x) \geq v \text{ for all } x \in S \quad (20)$$

$$\sum_j \mu_j d(S_i, x^j) \leq v \text{ for all } i = 1, \dots, m. \quad (21)$$

Since $d(S_i, x)$ are convex functions, the second inequality (1) implies

$$d\left(S_i, \sum_j \mu_j x^j\right) \leq v. \quad (22)$$

The game admits a pure optimal $x^\circ = \sum_j \mu_j x^j$ for Player II. We are also given

$$\bigcap_{i \neq j} S_i \neq \emptyset, \quad j = 1, 2, \dots, m.$$

For any optimal mixed strategy $\xi = (\xi_1, \xi_2, \dots, \xi_m)$ of Player I, if any $\xi_i = 0$ and we choose an $x^* \in \bigcap_{i \neq 1} S_i$, then from (20) we have $0 \geq v$ and thus $v = 0$. When the value v is zero, the third inequality (22) shows that $x^\circ \in \bigcap_i S_i$. If $\xi > 0$, the second inequality (21) will become an equality for all i and we have $d(S_i, x^\circ) = v$. But for $x^\circ \in S$, we have $v = 0$ and $x^\circ \in \bigcap_{i=1}^m S_i$. (See Raghavan (1973) for other applications.)

Ky Fan-Sion Minimax Theorems

General minimax theorems are concerned with the following problem: Given two arbitrary sets X , Y and a real function $K : X \times Y \rightarrow \mathbf{R}$, under what conditions on K , X , Y can one assert

$$\sup_{x \in X} \inf_{y \in Y} K(x, y) = \inf_{y \in Y} \sup_{x \in X} K(x, y).$$

A standard technique for proving general minimax theorems is to reduce the problem to the minimax theorem for matrix games. Such a reduction is often possible with some form of compactness of the space X or Y and a suitable continuity and convexity or quasi-convexity of the kernel K .

Definition 28 A function $f : X \rightarrow \mathbf{R}$ is upper-semi-continuous on X if and only if for any real c ,

$\{x : f(x) < c\}$ is open in X . A function $f: X \rightarrow \mathbf{R}$ is lower semi-continuous in X if and only if for any real c , $\{x : f(x) > c\}$ is open in X .

Definition 29 Let X be a convex subset of a topological vector space. A function $f: X \rightarrow \mathbf{R}$ is quasi-convex if and only if for each real c , the set $\{x : f(x) < c\}$ is convex. A function g is quasi-concave if and only if $-g$ is quasi-convex. Clearly any convex function (concave function) is quasi-convex (conceive).

The following minimax theorem is a special case of more general minimax theorems due to Ky Fan (1953) and Sion (1958).

Theorem 30 Let X, Y be compact convex subsets of linear topological spaces. Let $K: X \times Y \rightarrow \mathbf{R}$ be upper semi continuous (u.s.c) in x (for each fixed y) and lower semi continuous (l.s.c) in y (for each x). Let $K(x, y)$ be quasi-concave in x and quasi-convex in y . Then

$$\max_{x \in X} \min_{y \in Y} K(x, y) = \min_{y \in Y} \max_{x \in X} K(x, y).$$

Proof The compactness of spaces and the u. s. c, l. s. c conditions guarantee the existence of $\max_{x \in X} \min_{y \in Y} K(x, y)$ and $\min_{y \in Y} \max_{x \in X} K(x, y)$. We always have

$$\max_{x \in X} \min_{y \in Y} K(x, y) \leq \min_{y \in Y} \max_{x \in X} K(x, y).$$

If possible let $\max_{x \in X} \min_{y \in Y} K(x, y) < c < \min_{y \in Y} \max_{x \in X} K(x, y)$. Let $A_x = \{y : K(x, y) < c\}$ and $B_y = \{x : K(x, y) > c\}$. Therefore we have finite subsets $X_1 \subseteq X, Y_1 \subseteq Y$ such that for each $y \in Y$ and hence for each $y \in \text{Con } Y_1$, there is an $x \in X_1$ with $K(x, y) > c$ and for each $x \in X$ and hence for each $x \in \text{Con } X_1$, there is a $y \in Y_1$, with $K(x, y) < c$. Without loss of generality let the finite sets X_1, Y_1 be with *minimum* cardinality m and n satisfying the above conditions. The minimum cardinality conditions have the following implications. The sets $S_i = \{y : K(x_i, y) \leq c\} \cap \text{Con } Y_1$ are non-empty and convex.

Further $\bigcap_{i \neq j} S_i \neq \emptyset$ for all $j = 1, \dots, n$, but $\bigcap_{i=1}^n S_i = \emptyset$. Now by Berge's theorem (Subsect. "Geometric Consequences"), the union of the sets

S_i cannot be convex. Therefore there exists $y_0 \in \text{Con } Y_1$, with $K(x, y_0) > c$ for all $x \in X_1$. Since $K(., y_0)$ is quasi-concave we have $K(x, y_0) > c$ for all $x \in \text{Con } X_1$. Similarly there exists an $x_0 \in \text{Con } X_1$ such that $K(x_0, y) < c$ for all $y \in Y_1$ and hence for all $y \in \text{Con } Y_1$ (by quasi-convexity of $K(x_0, y)$). Hence $c < K(x_0, y_0) < c$, and we have a contradiction.

When the sets X, Y are mixed strategies (probability measures) for suitable Borel spaces one can find value for some games with optimal strategies for one player but not for both. See (Parthasarathy and Raghavan 1971); Alpern and Gal (2003).

Applications of Infinite Games

S-games and Discriminant Analysis

Motivated by Fisher's enquiries (Fisher 1936) into the problem of classifying a randomly observed human skull into one of several known populations, discriminant analysis has exploded into a major statistical tool with applications to diverse problems in business, social sciences and biological sciences (Johnson and Wichern 2007; Rao 1952).

Example 31 A population Π has a probability density which is either f_1 or f_2 where f_i is multi-variate normal with mean vector $\mu_i, i = 1, 2$ and variance covariance matrix, Σ , the same for both f_1 and f_2 . Given an observation X from population Π the problem is to classify the observation into the proper population with density f_1 or f_2 . The costs of misclassifications are $c(1/2) > 0$ and $c(2/1) > 0$ where $c(i/j)$ is the cost of misclassifying an observation from Π_j to Π_i . The aim of the statistician is to find a suitable decision procedure that minimizes the worst risk possible.

This can be treated as an S-game where the pure strategy for Player I (nature) is the secret choice of the population and a pure strategy for the statistician (Player II) is to partition the sample space into two disjoint sets (T_1, T_2) such that observations falling in T_1 are classified as from Π_1 and observations falling in T_2 are classified as

from Π_2 . The payoffs to Player I (Nature) when the observation is chosen from Π_1 , Π_2 is given by the risks (expected costs):

$$\begin{aligned} r(1, (T_1, T_2)) &= c(2/1) \int_{T_2} f_1(x) dx, \\ r(2, (T_1, T_2)) &= c(1/2) \int_{T_1} f_2(x) dx. \end{aligned}$$

The following theorem, based on an extension of Lyapunov's theorem for non-atomic vector measures (Lindenstrauss 1966), is due to Blackwell (1951), Dvoretzky-Wald and Wolfowitz (1951).

Theorem 32 Let

$$S = \left\{ (s_1, s_2) : s_1 = c(2/1) \int_{T_2} f_1(x) dx, \right. \\ \left. s_2 = c(1/2) \int_{T_1} f_2(x) dx, (T_1, T_2) \in \mathcal{T} \right\}$$

where $\mathcal{T}$ is the collection of all measurable partitions of the sample space. Then S is a closed bounded convex set.

We know from the theorem on S games (Theorem 22) that Player II (the statistician) has a minimax strategy which is pure. If v is the value of the game and if (ξ_1^*, ξ_2^*) is an optimal strategy for Player I then we have:

$$\xi_1^* \cdot c(2/1) \int_{T_2} f_1(x) dx + \xi_2^* \cdot c(1/2) \int_{T_1} f_2(x) dx \geq v$$

for all measurable partitions $\mathcal{T}$. For any general partition $\mathcal{T}$, the above expected payoff to I simplifies to:

$$\xi_1^* c(2/1) + \int_{T_1} [\xi_2^* \cdot c(1/2) f_2(x) - \xi_1^* c(2/1)] f_1(x) dx.$$

It is minimized whenever the integrand is ≤ 0 on T_1 . Thus the optimal pure strategy (T_1^*, T_2^*) satisfies:

$$T_1^* = \{x : \xi_2^* c(1/2) f_2(x) - \xi_1^* c(2/1) f_1(x) \leq 0\}$$

This is equivalent to

$$T_1^* = \left\{ x : U(x) = (\mu_1 - \mu_2)^T \sum^{-1} x - \frac{1}{2} (\mu_1 - \mu_2)^T \sum^{-1} (\mu_1 + \mu_2) \geq k \right\}$$

and

$$T_2^* = \left\{ U(x) = (\mu_1 - \mu_2)^T \sum^{-1} x - \frac{1}{2} (\mu_1 - \mu_2)^T \sum^{-1} (\mu_1 + \mu_2) < k \right\}$$

for some suitable k . Let $\alpha = (\mu_1 - \mu_2)^T \sum^{-1} (\mu_1 - \mu_2)$.

The random variable U is univariate normal with mean $\frac{\alpha}{2}$ and variance α if $x \in \Pi_1$. The random variable U has mean $\frac{-\alpha}{2}$ and variance α if $x \in \Pi_2$. The minimax strategy for the statistician will be such that

$$\begin{aligned} c(2/1) \int_{-\infty}^{\frac{1}{\sqrt{\alpha}}(k-\frac{\alpha}{2})} \frac{1}{\sqrt{2\pi}} e^{-\frac{y^2}{2}} dy \\ = c(1/2) \int_{\frac{1}{\sqrt{\alpha}}(k+\frac{\alpha}{2})}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{y^2}{2}} dy. \end{aligned}$$

The value of k can be determined by trial and error.

General Minimax Theorem and Statistical Estimation

Example 33 A coin falls heads with unknown probability θ . Not knowing the true value of θ a statistician wants to estimate θ based on a single toss of the coin with squared error loss.

Of course, the outcome is either heads or tails. If heads he can estimate $\hat{\theta}$ as x and if the outcome is tails, he can estimate $\hat{\theta}$ as y . To keep the problem zero-sum, the statistician pays a penalty $(\theta - x)^2$ when he proposes x as the estimate and pays a penalty $(\theta - y)^2$ when he proposes y as the estimate. Thus the *expected loss or risk* to the statistician is given by

$$\theta(\theta - x)^2 + (1 - \theta)(\theta - y)^2.$$

It was Abraham Wald (1950) who pioneered this game theoretic approach. The problem of estimation is one of discovering the true state of nature based on partial knowledge about nature revealed by experiments. While nature reveals in part, the aim of nature is to conceal its secrets. The statistician has to make his decisions by using observed information about nature. We can think of this as an ordinary zero sum two person game where the pure strategy for nature is to choose any $\theta \in [0, 1]$ and a pure strategy for Player II (statistician) is any point in the unit square $I = \{(x, y) : 0 \leq x \leq 1, 0 \leq y \leq 1\}$ with the payoff given above. Expanding the payoff function we get

$$K(\theta, (x, y)) = \theta^2(-2x + 1 + 2y) + \theta(x^2 - 2y - y^2) + y^2.$$

The statistician may try to choose his strategy in such a way that no matter what θ is chosen by nature, it has no effect on his penalty given the choice he makes for x, y . We call such a strategy an *equalizer*. We have an equalizer strategy if we can make the above payoff independent of θ . In fact we have used this trick earlier while dealing with behavior strategies.

For example by choosing x and y such that the coefficient of θ^2 and θ are zero we can make the payoff expression independent of θ . We find that $x = \frac{3}{4}, y = \frac{1}{4}$ is the solution. While this may guarantee the statistician a constant risk of $\frac{1}{16}$, one may wonder whether this is the best? We will show that it is best by finding a mixed strategy for mother nature that guarantees an expected payoff of $\frac{1}{16}$ no matter what (x, y) is proposed by the statistician. Let $F(\theta)$ be the cumulative distribution function that is optimal for mother nature. Integrating the above payoff with respect to F we get

$$K(F, (x, y)) = (-2x + 1 + 2y) \int_0^1 \theta^2 dF(\theta) + (x^2 - 2y - y^2) \int_0^1 \theta dF(\theta) + y^2.$$

Let

$$m_1 = \int_0^1 \theta dF(\theta) \text{ and } m_2 = \int_0^1 \theta^2 dF(\theta).$$

In terms of m_1, m_2, x, y the expected payoff can be rewritten as

$$L((m_1, m_2), (x, y)) = m_2(-2x + 1 + 2y) + m_1(x^2 - 2y - y^2) + y^2.$$

For fixed values of m_2, m_1 the minimum value of $m_2(-2x + 1 + 2y) + m_1(x^2 - 2y - y^2) + y^2$ must satisfy the first order conditions

$$\frac{m_2}{m_1} = x^*, \text{ and } \frac{m_2 - m_1}{m_1 - 1} = y^*.$$

If $m_1 = 1/2$ and $m_2 = 3/8$ then $x^* = 3/4$ and $y^* = 1/4$ is optimal. In fact, a simple probability distribution can be found by choosing the point $1/2 - \alpha$ with probability $1/2$ and $1/2 + \alpha$ with probability $1/2$. Such a distribution will have mean $m_1 = 1/2$ and second moment $m_2 = (1/2)(1/2 - \alpha)^2 + 1/2(1/2 + \alpha)^2 = 1/4 + \alpha^2$. When $m_2 = 3/8$, we get $\alpha^2 = 1/8$. Thus the optimal strategy for nature also called the *least favorable distribution*, is to toss an ordinary unbiased coin and if it is heads, select a coin which falls heads with probability $(\sqrt{2} - 1)/(2\sqrt{2})$ and if the ordinary coin toss is tails, select a coin which falls heads with probability $(\sqrt{2} + 1)/(2\sqrt{2})$. For further applications of zero-sum two person games to statistical decision theory see Ferguson (1967).

In general it is difficult to characterize the nature of optimal mixed strategies even for a C^∞ payoff function. See (Karlin 1959). Another rich source for infinite games do appear in many simplified parlor games. We give a simplified poker model due to Borel (See the reference under Ville (1938)).

Borel's Poker Model

Two risk neutral players I and II initiate a game by adding an ante of \$1 each to a pot. Players I and II are dealt a *hand*, a random value u of U and a random value v of V respectively. Here U, V are

independent random variables uniformly distributed on $[0, 1]$.

The game begins with Player I. After seeing his hand he either *folds* losing the pot to Player II, or *raises* by adding \$1 to the pot. When Player I raises, Player II after seeing his hand can either *fold*, losing the pot to Player I, or *call* by adding \$1 to the pot. If Player II calls, the game ends and the player with the better hand wins the entire pot. Suppose players I and II restrict themselves to using only strategies g, h respectively where

$$\begin{aligned} \text{Player I : } g(u) &= \begin{cases} \text{for if } u \leq x \\ \text{raise if } u > x, \end{cases} \\ \text{Player II : } h(v) &= \begin{cases} \text{for if } v \leq y \\ \text{raise if } v > y. \end{cases} \end{aligned}$$

The computational details of the expected payoff $K(x, y)$ based on the above partitions of the u, v space is given below in Tables 3 and 4.

The payoff $K(x, y)$ is simply the sum of each area times the local payoff given by

$$K(x, y) = 2y^2 - 3xy + x - y \text{ when } x < y.$$

The payoff $K(x, y)$ is simply the sum of each area times the local payoff given by


$$K(x, y) = \begin{cases} -2x^2 + xy + x - y & \text{for } x > y \\ 2y^2 - 3xy + x - y & \text{for } x < y. \end{cases}$$

Also $K(x, y)$ is continuous on the diagonal and concave in x and convex in y . By the general minimax theorem of Ky Fan or Sion there exist x^*, y^* pure optimal strategies for $K(x, y)$. We will find the value of the game by explicitly computing $\min_y K(x, y)$ for each x and then taking the maximum over x .


Let $0 < \bar{x} < 1$. (Intuitively we see that in our search for saddle point, $x = 0$ or $x = 1$ are quite unlikely.)

$\min_{0 \leq y \leq 1} K(\bar{x}, y) = \min \{ \inf_{y < \bar{x}} K(\bar{x}, y), \min_{y \geq \bar{x}} K(\bar{x}, y) \}$. Observe that $\inf_{y < \bar{x}} K(\bar{x}, y) = \min(-2\bar{x}^2 + \bar{x}y + \bar{x} - y) = y(\bar{x} - 1)$ terms not involving y .

Clearly the infimum is attained at $y = \bar{x}$ giving a value $-\bar{x}^2$. Now for $y > \bar{x}$ we have $K(\bar{x}, y) = 2y^2 - 3\bar{x}y + \bar{x} - y$. This being a convex function in y , it has its minimum either at the boundary or at an interior point of the interval $[\bar{x}, 1]$. We have $K(\bar{x}, \bar{x}) = -\bar{x}^2$, $K(\bar{x}, 1) = 1 - 2\bar{x}$ or the minimum is attained at $y = (3\bar{x} + 1)/4$, the unique solution to $(\partial K(\bar{x}, y))/(\partial y) = 0$. At $y = (3\bar{x} + 1)/4$, the value of $K(\bar{x}, y) = (-9\bar{x}^2 + 2\bar{x} - 1)/8$. Clearly for $\bar{x} \neq 1$ we have $(-9\bar{x}^2 + 2\bar{x} - 1)/8 < -\bar{x}^2 < 1 - 2\bar{x}$. Thus $\min_y K(x, y) = (-9x^2 + 2x - 1)/8$. Now $\max_x \min_y K(x, y) = \max_x (-9x^2 + 2x - 1)/8 = -1/9$. It is attained at $x^* = 1/9$. Now $\min_y K(1/9, y)$ is attained at $y^* = 1/3$.



Zero-Sum Two Person Games, Fig. 5 Poker partition when $x < y$



Zero-Sum Two Person Games, Fig. 6 Poker partition when $x > y$

Zero-Sum Two Person Games, Table 3

Outcome	Action by players:	Region	Area	Payoff
$u \leq x$	I drops out	$A \cup D \cup F$	x	-1
$u > x, v \leq y$	II drops out	$E \cup G$	$y(1-x)$	1
$u < v, u > x, v > y$	both raise	B	$\frac{(1-y)(1-2x+y)}{2}$	-2
$u > v, u > x, v > y$	both raise	C	$\frac{1}{2}(1-y)^2$	2

Zero-Sum Two Person Games, Table 4

Outcome	Action by players:	Region	Area	Payoff
$u \leq x$	I drops out	$H \cup I \cup J \cup K$	x	-1
$u > x, v \leq y$	II drops out	N	$y(1-x)$	1
$u < v, u > x, v > y$	both raise	L	$\frac{(1-x)^2}{2}$	-2
$u > v, u > x, v > y$	both raise	M	$\frac{1}{2}(1-x)(1+x-2y)$	2

Thus the value of the game is $-1/9$. A good pure strategy for I is to raise when $x > 1/9$ and a good pure strategy for II is to call when $y > 1/3$.

For other poker models see von Neumann and Morgenstern (1947), Blackwell and Bellman (1949), Binmore (1992) and Ferguson and Ferguson (2003) who discuss other discrete and continuous poker models.

War Duels and Discontinuous Payoffs on the Unit Square

While general minimax theorems make some stringent continuity assumptions, rarely they are satisfied in modeling many war duels as games on the unit square. The payoffs are often discontinuous along the diagonal. The following is an example of this kind.

Example 34 Players I and II start at a distance = 1 from a balloon, and walk toward the balloon at the

same speed. Each player carries a noisy gun which has been loaded with just one bullet. Player I's accuracy at x is $p(x)$ and Player II's accuracy at x is $q(x)$ where x is the distance traveled from the starting point. Because the guns can be heard, each player knows whether or not the other player has fired. The player who fires and hits the balloon first wins the game.

Some natural assumptions are

- The functions p, q are continuous and strictly increasing
- $p(0) = q(0) = 0$ and $p(1) = q(1) = 1$.

If a player fires and misses then the opponent can wait until his accuracy = 1 and then fire. Player I decides to shoot after traveling distance x provided the opponent has not yet used his bullet. Player II decides to shoot after traveling distance y provided the opponent has not yet used his bullet. The following payoff reflects the outcome of this strategy choice.

$$K(x, y) = \begin{cases} (1)p(x) + (-1)(1 - p(x)) = 2p(x) - 1 & \text{when } x < y \\ p(x) - q(x) & \text{when } x = y, \\ (-1)q(y) + (1 - q(y))(1) = 1 - 2q(y) & \text{when } x > y. \end{cases}$$

We claim that the players have optimal pure strategies and there is a saddle point. Consider $\min_y K(x, y) = \min \{2p(x) - 1, p(x) - q(x), \inf_{y < x} (1 - 2q(y))\}$. We can replace $\inf_{y < x} (1 - 2q(y))$ by $1 - 2q(x)$. Thus we get

$$\min_{0 \leq y \leq 1} K(x, y) = \min \{2p(x) - 1, p(x) - q(x), (1 - 2q(x))\}.$$

Since the middle function $p(x) - q(x)$ is the average of the two functions $2p(x) - 1$ and $1 - 2q(x)$, the minimum of the three functions is simply the $\min \{2p(x) - 1, (1 - 2q(x))\}$. While the first function is increasing in x the second one is decreasing in x . Thus the $\max_x \min_y K(x, y)$ is the solution to the eq. $2p(x) - 1 = 1 - 2q(x)$. There is a unique solution to the equation $p(x) + q(x) = 1$ as both functions are strictly increasing. Let x^* be that solution point. We get $p(x^*) + q(x^*) = 1$. The value v satisfies

$$\begin{aligned} v &= p(x^*) - q(x^*) \\ &= \min_y \max_x K(x, y) \\ &= \max_x \min_y K(x, y). \end{aligned}$$

Remark 35 In the above example the winner's

payoff is 1 whether player I wins or player II wins. This is crucial for the existence of optimal pure strategies for the two players. Suppose the payoff has the following structure:

- α if I wins
- β if II wins
- γ if neither wins
- 0 if both shoot accurately and ends in a draw.

Depending on these values while the value exists only one player may have optimal pure strategy while the other may have only an *epsilon* optimal pure strategy, close to but not the same as the solution to the equation $p(x^*) + q(x^*) = 1$.

Example 36 (silent duel) Players I and II have the same accuracy $p(x) = q(x) = x$. However, in this duel, the players are both deaf so they do not know whether or not the opponent has fired. The payoff function is given by

$$K(x, y) = \begin{cases} (1)x + (-1)(y)(1 - x) = x - y + xy & \text{when } x < y, \\ 0 & \text{when } x = y, \\ (-1)y + (1)(1 - y)x = x - y - xy & \text{when } x > y. \end{cases}$$

This game has no saddle point. In this case, the value of the game if it exists must be zero. One can directly verify that the density

$$f(t) = \begin{cases} 0 & 0 \leq t < 1/3 \\ \frac{1}{4}t^{-3} & 1/3 \leq t \leq 1. \end{cases}$$

is optimal for both players with value zero.

Remark 37 In the above game suppose exactly one player, say Player II is deaf. Treating winners symmetrically, we can represent this game with the payoff

$$K(\xi, \eta) = \begin{cases} (1)\xi - (1 - \xi)\eta & \xi < \eta \\ 0 & \xi = \eta \\ (-1)\eta + (1 - \eta)(1) & \xi > \eta. \end{cases}$$

Let $a = \sqrt{6} - 2$. Then the game has value $v = 1 - 2a = 5 - 2\sqrt{6}$. An optimal strategy for Player I is given by the density

$$f(\xi) = \begin{cases} 0 & \text{for } 0 \leq \xi < a \\ \sqrt{2}a(\xi^2 + 2\xi - 1)^{-\frac{3}{2}} & \text{for } a \leq \xi \leq 1. \end{cases}$$

For Player II it is given by

$$g(\eta) = \begin{cases} 0 & \text{for } 0 \leq \xi < a \\ 2\sqrt{2} \cdot \frac{a}{2+a} (\eta^2 + 2\eta - 1) & \sqrt{6} - 2 \leq \eta \leq 1 \end{cases}$$

with an additional mass of $\frac{a}{2+a}$ at $\eta = 1$. The deaf player has to maintain a sizeable suspicion until the very end!

The study of war duels is intertwined with the study of positive solutions to integral equations. The theorems of Krein and Rutman (1950) on positive operators and their positive eigenfunctions are central to this analysis. For further details see (Dresher 1962; Karlin 1959; Radzik 1988).

Dynamic versions of zero sum two person games where the players move among several games according to some Markovian transition law leads to the theory of stochastic games (Shapley 1953). The study of value and the complex structure of optimal strategies of zero-sum two person stochastic games is an active area of research with applications to many dynamic models of competition (Filar and Vrieze 1996).

While stochastic games study movement in discrete time, the parallel development of dynamic games in continuous time was initiated by Isaacs in several Rand reports culminating in his monograph on Differential games (Isaacs 1965). Many military problems of pursuit and evasion, attrition and attack lead to games where the trajectory of a moving object is being steered continuously by the actions of two players. Introducing several natural examples Isaacs explicitly tried to solve many of these games via some heuristic principles. The minimax version of Bellman's optimality principle lead to the so called Isaacs Bellman equations. This highly non-linear partial differential equation on the value function plays a key role in this study.

Epilogue

The rich theory of zero sum two person games that we have discussed so far hinges on the fundamental notions of value and optimal strategies. When either zero sum or two person assumption is dropped, the games cease to have such well defined notions with independent standing. In trying to extend the notion of optimal strategies and the minimax theorem for bimatrix games Nash (1950) introduced the concept of an equilibrium

point. Even more than this extension it is this concept which is simply the most seminal solution concept for non-cooperative game theory. This Nash equilibrium solution is extendable to any non-cooperative N person game in both extensive and normal form. Many of the local properties of optimal strategies and the proof techniques of zero sum two person games do play a significant role in understanding Nash equilibria and their structure (Bubelis 1979; Jansen 1981; Kreps 1974; Raghavan 1970, 1973).

When a non-zero sum game is played repeatedly, the players can peg their future actions on past history. This leads to a rich theory of equilibria for repeated games (Sorin 1992). Some of them, like the repeated game model of Prisoner's dilemma impose tacitly, actual cooperation at equilibrium among otherwise non-cooperative players. This was first recognized by Schelling (1960), and was later formalized by the so called folk theorem for repeated games (Aumann and Shapley 1986; Axelrod and Hamilton 1981). Indeed Nash equilibrium per se is a weak solution concept is also one of the main messages of Folk theorem for repeated games. With a plethora of equilibria, it fails to have an appeal without further refinements. It was Selten who initiated the need for refining equilibria and came up with the notion of subgame perfect equilibria (Selten 1975). It turns out that subgame perfect equilibria are the natural solutions for many problems in sequential bargaining (Rubinstein 1982). Often one searches for specific types of equilibria like symmetric equilibria, or Bayesian equilibria for games with incomplete information (Harsanyi 1967). Auctions as games exhibit this diversity of equilibria and Harsanyi's Bayesian equilibria turn out to be the most appropriate solution concept for this class of games (Myerson 1991).

Zero sum two person games and their solutions will continue to inspire researchers in all aspects non-cooperative game theory.

Acknowledgment The author wishes to acknowledge the unknown referee's detailed comments in the revision of the first draft. More importantly he drew the author's attention to the topic of search games and other combinatorial games. The author would like to thank Ms. Patricia Collins for her assistance in her detailed editing of the first draft of

this manuscript. The author owes special thanks to Mr. Ramanujan Raghavan and Dr. A.V. Lakshmi Narayanan for their help in incorporating the graphics drawings into the latex file.

Bibliography

- Alpern S, Gal S (2003) The theory of search games and rendezvous. Springer
- Aumann RJ (1981) Survey of repeated games, essays in game theory and mathematical economics. In: Honor of Oscar Morgenstern. Bibliographisches Institut, Mannheim, pp 11–42
- Axelrod R, Hamilton WD (1981) The evolution of cooperation. *Science* 211:1390–1396
- Bapat RB, Raghavan TES (1997) Nonnegative matrices and applications. In: *Encyclopedia in mathematics*. Cambridge University Press, Cambridge
- Bellman R, Blackwell D (1949) Some two person games involving bluffing. *Proc Natl Acad Sci U S A* 35:600–605
- Berge C (1963) Topological spaces. Oliver Boyd, Edinburgh
- Berger U (2007) Brown's original fictitious play. *J Econ Theory* 135:572–578
- Berlekamp ER, Conway JH, Guy RK (1982) Winning ways for your mathematical plays, vol 1, 2. Academic, New York
- Binmore K (1992) Fun and game theory a text on game theory. Lexington, DC Heath
- Blackwell D (1951) On a theorem of Lyapunov. *Ann Math Stat* 22:112–114
- Blackwell D (1961) Minimax and irreducible matrices. *Math J Anal Appl* 3:37–39
- Blackwell D, Girshick GA (1954) Theory of games and statistical decisions. Wiley, New York
- Bouton CL (1902) Nim-a game with a complete mathematical theory. *Ann Math* 3(2):35–39
- Brown GW (1951) Iterative solution of games by fictitious play. In: Koopmans TC (ed) Activity analysis of production and allocation. Wiley, New York, pp 374–376
- Bubelis V (1979) On equilibria in finite games. *Int J Game Theory* 8:65–79
- Chin H, Parthasarathy T, Raghavan TES (1973) Structure of equilibria in N -person noncooperative games. *Int Game J Theory* 3:1–19
- Conway JH (1982) On numbers and games, monograph 16. London Mathematical Society, London
- Dantzig GB (1951) A proof of the equivalence of the programming problem and the game problem. In: Koopman's activity analysis of production and allocation, Cowles Conumeson monograph 13. Wiley, New York, pp 333–335
- Dresher M (1962) Games of strategy. Prentice Hall, Englewood Cliffs
- Dvoretzky A, Wald A, Wolfowitz J (1951) Elimination of randomization in certain statistical decision problems and zero-sum two-person games. *Ann Math Stat* 22:1–21
- Fan K (1953) Minimax theorems. *Proc Natl Acad Sci Wash* 39:42–47
- Ferguson TS (1967) Mathematical stat, a decision theoretic approach. Academic, New York
- Ferguson C, Ferguson TS (2003) On the Borel and von Neumann poker models. *Game Theory Appl* 9:17–32
- Filar JA, Vrieze OJ (1996) Competitive Markov decision processes. Springer, Berlin
- Fisher RA (1936) The use of multiple measurements in taxonomic problems. *Ann Eugenics* 7:179–188
- Fourier JB (1890) Second extrait. In: Darboux GO (ed) Gauthiers Villars, Paris, pp 325–328; English Translation: Kohler DA (1973)
- Gal S (1980) Search games. Academic, New York
- Gale D (1979) The game of Hex and the Brouwer fixed-point theorem. *Am Math Mon* 86:818–827
- Gale D, Kuhn HW, Tucker AW (1951) Linear programming and the theory of games. In: Activity analysis of production and allocation. Wiley, New York, pp 317–329
- Harsanyi JC (1967) Games with incomplete information played by Bayesian players, parts I, II, and III. *Sci Manag* 14:159–182; 32–334; 486–502
- Isaacs R (1965) Differential games: mathematical a theory with applications to warfare and pursuit. Control and optimization. Wiley, New York; Dover Paperback Edition, 1999
- Jansen MJM (1981) Regularity and stability of equilibrium points of bimatrix games. *Math Oper Res* 6:530–550
- Johnson RA, Wichern DW (2007) Applied multivariate statistical analysis, 6th edn. Prentice Hall, New York
- Kakutani S (1941) A generalization of Brouwer's fixed point theorem. *Duke Math J* 8:457–459
- Kantorowich LV (1960) Mathematical methods of organizing and planning production. *Manag Sci* 7:366–422
- Kaplansky I (1945) A contribution to von Neumann's theory of games. *Ann Math* 46:474–479
- Karlin S (1959) Mathematical methods and theory in games, programming and econs, vol 1, 2. Addison Wesley, New York
- Kohler DA (1973) Translation of a report by Fourier on his work on linear inequalities. *Opsearch* 10:38–42
- Krein MG, Rutmann MA (1950) Linear operators leaving invariant a cone in a Banach space. *Amer Math Soc Transl* 26:1–128
- Kreps VL (1974) Bimatrix games with unique equilibrium points. *Int Game J Theory* 3:115–118
- Krishna V, Sjostrom T (1998) On the convergence of fictitious play. *Math Oper Res* 23:479–511
- Kuhn HW (1953) Extensive games and the problem of information. Contributions to the theory of games. *Ann Math Stud* 28:193–216
- Lindenstrauss J (1966) A short proof of Liapounoff's convexity theorem. *Math J Mech* 15:971–972
- Loomis IH (1946) On a theorem of von Neumann. *Proc Nat Acad Sci Wash* 32:213–215
- Miyazawa K (1961) On the convergence of the learning process in a 2×2 non-zero-sum two-person game. Econometric research program, research memorandum no. 33. Princeton University, Princeton

- Monderer D, Shapley LS (1996) Potential LS games. *Games Econ Behav* 14:124–143
- Myerson R (1991) *Game theory. Analysis of conflict*. Harvard University Press, Cambridge, MA
- Nash JF (1950) Equilibrium points in n -person games. *Proc Natl Acad Sci Wash* 88:48–49
- Owen G (1985) *Game theory*, 2nd edn. Academic, New York
- Parthasarathy T, Raghavan TES (1971) *Some topics in two person games*. Elsevier, New York
- Radzik T (1988) Games of timing related to distribution of resources. *Optim J Theory Appl* 58(443–471):473–500
- Raghavan TES (1970) Completely mixed strategies in bimatrix games. *Lond J Math Soc* 2:709–712
- Raghavan TES (1973) Some geometric consequences of a game theoretic result. *Math J Anal Appl* 43:26–30
- Rao CR (1952) *Advanced statistical methods in biometric research*. Wiley, New York
- Reijnierse JH, Potters JAM (1993) Search games with immobile hider. *Int J Game Theory* 21:385–394
- Robinson J (1951) An iterative method of solving a game. *Ann Math* 54:296–301
- Rubinstein A (1982) Perfect equilibrium in a bargaining model. *Econometrica* 50:97–109
- Schelling TC (1960) *The strategy of conflict*. Harvard University Press, Cambridge, MA
- Selten R (1975) Reexamination of the perfectness concept for equilibrium points in extensive games. *Int Game J Theory* 4:25–55
- Shapley LS (1953) Stochastic games. *Proc Natl Acad of Sci USA* 39:1095–1100
- Sion M (1958) On general minimax theorem. *Pac J Math* 8:171–176
- Sorin S (1992) Repeated games with complete information, ster 4. In: Aumann RJ, Hart S (eds) *Handbook of game theory*, vol 1. North Holland, Amsterdam, pp 71–103
- Thuijsman F, Raghavan TES (1997) Stochastic games with switching control or ARAT structure. *Int Game J Theory* 26:403–408
- Ville J (1938) Note sur la theorie generale des jeux ou intervient l'habilité des joueurs. In: Borel E, Ville J (eds) *Applications aux jeux de hasard*, tome IV, fascicule II of the *Traite du calcul des probabilités et de ses applications*, by Emile Borel
- von Neumann J (1928) Zur Theorie der Gesellschaftspiele. *Math Ann* 100:295–320
- von Neumann J, Morgenstern O (1947) *Theory of games and economic behavior*, 2nd edn. Princeton University Press, Princeton
- von Stengel B (1996) Efficient computation of behavior strategies. *Games Econ Behav* 14:220–246
- Wald A (1950) *Statistical decision functions*. Wiley, New York
- Weyl H (1950) Elementary proof of a minimax theorem due to von Neumann. *Ann Math Stud* 24:19–25
- Zermelo E (1913) Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels. In: *Proceedings of the fifth congress mathematicians*. Cambridge University Press, Cambridge, MA, pp 501–504

Stochastic Games

Yehuda John Levy¹ and Eilon Solan²

¹Adam Smith Business School, University of Glasgow, Glasgow, UK

²The School of Mathematical Sciences, Tel Aviv University, Tel Aviv, Israel

Article Outline

Glossary
Definition of the Subject and Its Importance
Strategies, Evaluations, and Equilibria
Zero-Sum Games
Multiplayer Games
Correlated Equilibrium
Imperfect Monitoring
Folk Theorems
Algorithms
Continuous-Time Games
Additional and Future Directions
Bibliography

Glossary

A correlated equilibrium An equilibrium in an extended game in which either at the outset of the game or at various points during the play of the game each player receives a private signal, and the vector of private signals is chosen according to a known joint probability distribution. In the extended game, a strategy of a player depends, in addition to past play, on the signals he received.

A stochastic game A repeated interaction between several participants in which the underlying state of the environment changes

We thank Eitan Altman, János Flesch, Yuval Heller, Jean-Jacques Herings, Ayala Mashiach-Yakovi, Andrzej Nowak, Ronald Peeters, T.E.S. Raghavan, Jérôme Renault, Nahum Shimkin, Robert Simon, Sylvain Sorin, William Sudderth, and Frank Thuijmsman, for their comments on an earlier version of the entry.

stochastically, and it depends on the decisions of the participants.

A strategy A rule that dictates how a participant in an interaction makes his decisions as a function of the observed behavior of the other participants and of the evolution of the environment.

An equilibrium A collection of strategies, one for each player, such that each player maximizes (or minimizes, in case of stage costs) his evaluation of stage payoffs given the strategies of the other players.

Evaluation of stage payoffs The way that a participant in an ongoing interaction evaluates the stream of stage payoffs that he receives (or stage costs that he pays) along the interaction.

Definition of the Subject and Its Importance

Stochastic games, first introduced by Shapley (1953), model dynamic interactions in which the environment changes in response to the behavior of the players. Formally, a stochastic game is a tuple $G = \langle N, S, (\mathcal{A}_i, A_i, u_i)_{i \in N, q} \rangle$ where:

- N is a set of players.
- S is a state space. If S is uncountable, it is supplemented with a σ -algebra of measurable sets.
- For every player $i \in N$, $\mathcal{A}_i$ is a set of actions for that player, and $A_i : S \rightarrow \mathcal{A}_i$ is a set-valued (measurable) function that assigns to each state $s \in S$ the set of actions $A_i(s)$ that are available to player i in state s . If $\mathcal{A}_i$ is uncountable, it is supplemented with an σ -algebra of measurable sets. Denote $SA = \{(s, a) : s \in S, a = (a_i)_{i \in N}, a_i \in A_i(s) \forall i \in N\}$. This is the set of all possible *action profiles*.
- For every player $i \in N$, $u_i : SA \rightarrow \mathbf{R}$ is a (measurable) *stage payoff function* for player i .
- $q : SA \rightarrow \Delta(S)$ is a (measurable) *transition function*, where $\Delta(S)$ is the space of probability distributions over S .

The game starts at an initial state $s^1 \in S$ and is played as follows. At each stage $t \in \mathbf{N}$, each player $i \in N$ chooses an action $a_i^t \in A_i(s^t)$ and receives the stage payoff $u_i(s^t, a^t)$, where $a^t = (a_i^t)_{i \in N}$, and the game moves to a new state s^{t+1} that is chosen according to the probability distribution $q(\cdot | s^t, a^t)$.

A few comments are in order:

1. A stochastic game lasts infinitely many stages. However, the model also captures finite interactions (of length t), by assuming the play moves, at stage t , to an absorbing state with payoff 0 to all players.
2. In particular, by setting $t = 1$, we see that stochastic games are a generalization of matrix games (games in normal form), which are played only once.
3. Stochastic games are also a generalization of repeated games, in which the players play the same matrix game over and over again. Indeed, a repeated game is equivalent to a stochastic game with a single state.
4. Stopping games are also a special case of stochastic games. In these games every player has two actions in all states, *continue* and *stop*. As long as all players choose *continue*, the stage payoff is 0; once at least one player chooses *stop*, the game moves to an absorbing state.
5. Markov decision problems (see, e.g., Puterman 1994) are stochastic games with a single player.
6. The transition function q governs the evolution of the game. It depends on the actions of all players and on the current state, so that all the players may influence the evolution of the game.
7. The payoff function u_i of player i depends on the current state as well as on the actions chosen by all players. Thus, a player's payoff depends not only on that player's choice but also on the behavior of the other players.
8. Though we refer to the functions $(u_i)_{i \in N}$ as "stage payoffs," with the implicit assumption that each player tries to maximize his payoff, in some applications these functions describe a stage *cost*, and then the implicit assumption is that each player tries to minimize his cost.
9. The action of a player at a given stage affects both his stage payoff and the evolution of the state variable, thereby affecting his future payoffs. These two sometimes contradicting effects make the optimization problem of the players quite intricate and the analysis of the game challenging.
10. The players receive a stage payoff at each stage. So far we did not mention how the players evaluate the infinite stream of stage payoffs that they receive, nor did we say what is their information at each stage: Do they observe the current state? Do they observe the actions of the other players? These issues will be discussed later.
11. The actions that are available to the players at each stage, the payoff functions, and the transition function all depend on the current state and not on past play (i.e., past states that the game visited and past actions that the players chose). This assumption is without loss of generality. Indeed, suppose that the actions available to the players at each stage, the payoff functions, and the transition function all depend on past play, as well as on the current state. For every $t \in \mathbf{N}$, let H_t be the set of all possible *histories* of length t , that is, all sequences of the form $(s^1, a^1, s^2, a^2, \dots, s^t)$, where $s^k \in S$ for every $k = 1, 2, 3, \dots, t$, $a^k = (a_i^k)_{i \in N}$ and a_i^k is an available action to player i at stage k , for every $k = 1, 2, \dots, t - 1$. Then the game is equivalent to a game with state space $H := \cup_{t \in \mathbf{N}} H_t$, in which the state variable captures past play and the state at stage t lies in H_t . In the new game, the sets of available actions, the payoff function, and the transition function depend on the current state rather than on all past play.

The interested reader is referred to Filar and Vrieze (1996), Sorin (2002), Neyman and Sorin (2003), and Başar and Zaccour (2017) for further reading on stochastic games. We now provide a few applications.

Example 1 Capital Accumulation (Levhari and Mirman 1980; Dutta and Sundaram 1992, 1993; Amir 1996; Nowak 2003c) *Two (or more) agents jointly own a natural resource or a productive asset; at every period they have to decide on the amount of the resource to consume.*

The amount that is not consumed grows by a known (or an unknown) fraction. Such a situation occurs, e.g., in fishery: fishermen from various countries fish in the same area, and each country sets a quota for its fishermen. Here the state variable is the current amount of resource, the action set is the amount of resource to be exploited in the current period, and the transition is influenced by the decisions of all the players, as well as possibly by the random growth of the resource.

Example 2 Taxation (Chari and Kehoe 1990; Phelan and Stacchetti 2001) A government sets a tax rate at every period. Each citizen decides at every period how much to work and, from the total amount of money he or she has, how much to consume; the rest is saved for the next period and grows by a known interest rate. Here the state is the amount of savings each citizen has; the stage payoff of a citizen depends on the amount of money that he consumed, on the amount of free time he has, and on the total amount of tax that the government collected. The stage payoff of the government may be the average stage payoff of the citizens, the amount of tax collected, or a mixture of the two.

Example 3 Communication Network (Sagduyu and Ephremides 2003) A single-cell system with one receiver and multiple uplink transmitters shares a single, slotted, synchronous classical collision channel. Assume that all transmitted packets have the same length and require one time unit, which is equal to one time slot, for transmission. Whenever a collision occurs, the users attempt to retransmit their packets in subsequent slots to resolve collision for reliable communication.

Here a state lists all relevant data for a given stage, e.g., the number of packets waiting at each transmitter, or the length of time each has been waiting to be transmitted. The players are the transmitters, and the action of each transmitter is which packet to transmit, if any. The stage cost may depend on the number of time slots that the transmitted packet waited, on the number of packets that have not been transmitted at that period, and possibly on additional variables.

The transition depends on the actions chosen by the players, but it has a stochastic component, which captures the number of new packets that arrive at the various transmitters during every time slot.

Example 4 Queues (Altman 2005) Individuals that require service have to choose whether to be served by a private slow service provider or by a powerful public service provider. This situation arises, e.g., when jobs can be executed on either a slow personal computer or a fast mainframe. Here a state lists the current load of the public and private service providers, and the cost is the time to be served.

The importance of stochastic games stems from the wide range of applications they encompass. Many repeated interactions can be recast as stochastic games; the vast array of theoretical results that have been obtained provide insights that can help in analyzing specific situations and suggesting proper behavior to the participants. In certain classes of games, algorithms that have been developed may be used to calculate such behavior.

Strategies, Evaluations, and Equilibria

So far we have not described the information that the players have at each stage. In most of the chapter, we assume that the players have complete information of past play; that is, at each stage t , they know the sequence $s^1, a^1, s^2, a^2, \dots, s^t$ of states that were visited in the past (including the current state) and the actions that were chosen by all players. This assumption is too strong for most applications, and in the sequel we will mention the consequences of its relaxation.

Since the players observe past play, a *pure strategy* for player i is a (measurable) function σ_i that assigns to every finite history $(s^1, a^1, s^2, a^2, \dots, s^t)$ an action $\sigma_i(s^1, a^1, s^2, a^2, \dots, s^t) \in A_i(s^t)$, with the interpretation that, at stage t , if the finite history $(s^1, a^1, s^2, a^2, \dots, s^t)$ occurred, player i plays the action $\sigma_i(s^1, a^1, s^2, a^2, \dots, s^t)$. If the player does not know the complete history, then a strategy for player i is a function that assigns to every possible information set, an

action that is available to the player when the player has this information. A *mixed strategy* for player i is a probability distribution over the set of his pure strategies. The space of mixed strategies of player i is denoted by Σ_i . A *behavior strategy* is a function σ_i that assigns to every finite history $(s^1, a^1, s^2, a^2, \dots, s^t)$ a mixed action $\sigma_i(s^1, a^1, s^2, a^2, \dots, s^t) \in \Delta(A_i(s^t))$, the set of probability distributions over $A_i(s^t)$. By Kuhn's theorem, if players have *perfect recall*, that is, no player forgets information that he knew in the past and each player remembers his own past choices, every behavior strategy is equivalent to some mixed strategy and vice versa.

A simple class of strategies is the class of *stationary strategies*; a strategy σ_i for player i is *stationary* if $\sigma_i(s^1, a^1, s^2, a^2, \dots, s^t)$ depends only on the current state s^t and not on past play $s^1, a^1, s^2, a^2, \dots, a^{t-1}$. A stationary strategy of player i can be identified with an element $x = (x_s)_{s \in S} \in \times_{s \in S} \Delta(A_i(s))$, with the interpretation that player i plays the mixed action x_s whenever the current state is s . Denote by $X_i = \times_{s \in S} \Delta(A_i(s))$ the space of stationary strategies of player i .

There are three common ways to evaluate the infinite stream of payoffs that the players receive in a stochastic game: the *finite-horizon evaluation*, in which a player considers the average payoff during the first T stages; the *discounted evaluation*, in which a player considers the discounted sum of his stage payoffs; and the *limsup evaluation*, in which a player considers the limsup of his long-run average payoffs. We now formally define these evaluations.

Every profile $\sigma = (\sigma_i)_{i \in N}$ of mixed strategies, together with the initial state, induces a probability distribution $\mathbf{P}_{s^1, \sigma}$ over the space of infinite plays $H_\infty := SA^N$. We denote the corresponding expectation operator by $\mathbf{E}_{s^1, \sigma}$.

Definition 5 Let σ be a profile of mixed strategies. For every finite horizon $T \in \mathbf{N}$, the T -stage payoff under σ for player i is

$$\gamma_i^T(s^1, \sigma) := \mathbf{E}_{s^1, \sigma} \left[\frac{1}{T} \sum_{t=1}^T u_i(s^t, a^t) \right].$$

For every discount rate $\lambda \in (0, 1]$, the λ -discounted payoff under σ for player i is

$$\gamma_i^\lambda(s^1, \sigma) := \mathbf{E}_{s^1, \sigma} \left[\lambda \sum_{t=1}^{\infty} (1 - \lambda)^{t-1} u_i(s^t, a^t) \right].$$

The limsup payoff under σ for player i is

$$\gamma_i^\infty(s^1, \sigma) := \mathbf{E}_{s^1, \sigma} \left[\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T u_i(s^t, a^t) \right].$$

The T -stage payoff captures the situation in which the interaction lasts exactly T stages. The λ -discounted evaluation captures the situation in which the game lasts “many” stages, and the player discounts stage payoffs – it is better to receive \$1 today than tomorrow. The limsup payoff also captures the situation in which the game lasts “many” stages, but here the player does not discount his payoffs, and the payoff at each given stage is insignificant as compared to the payoff in all other stages. Equivalently, one could consider the liminf payoff in which the player considers the liminf of the long-run average payoffs.

As usual, an equilibrium is a vector of strategies such that no player can profit by a unilateral deviation. For every player i and every strategy profile $\sigma = (\sigma_i)_{i \in N}$, we denote the strategy profile of all other players, except player i , by $\sigma_{-i} = (\sigma_j)_{j \neq i}$.

Definition 6 Let $\varepsilon \geq 0$. A profile of strategies σ is a T -stage ε -equilibrium if

$$\gamma_i^T(s^1, \sigma) \geq \gamma_i^T(s^1, \sigma'_i, \sigma_{-i}) - \varepsilon, \quad \forall s^1 \in S, \forall i \in N, \forall \sigma'_i \in \Sigma_i.$$

It is a λ -discounted ε -equilibrium if

$$\gamma_i^\lambda(s^1, \sigma) \geq \gamma_i^\lambda(s^1, \sigma'_i, \sigma_{-i}) - \varepsilon, \quad \forall s^1 \in S, \forall i \in N, \forall \sigma'_i \in \Sigma_i.$$

It is a limsup ε -equilibrium if

$$\gamma_i^\infty(s^1, \sigma) \geq \gamma_i^\infty(s^1, \sigma'_i, \sigma_{-i}) - \varepsilon, \quad \forall s^1 \in S, \forall i \in N, \forall \sigma'_i \in \Sigma_i.$$

The payoff that corresponds to an ε -equilibrium, that is, either one of the quantities $\gamma^T(s^1, \sigma)$, $\gamma^\lambda(s^1, \sigma)$, and $\gamma^\infty(s^1, \sigma)$, is called an ε -equilibrium payoff at the initial state s^1 .

A 0-equilibrium is termed an *equilibrium*. As we will see below, when both state and action spaces are finite, a T -stage and a λ -discounted equilibrium exist. However, when the state or action spaces are infinite, such an equilibrium may fail to exist, yet ε -equilibria may exist for every $\varepsilon > 0$.

As the length of the game T varies, or as the discount rate λ varies, the equilibrium strategy profile varies as well. A strategy profile that is an ε -equilibrium for every T sufficiently large, every

λ sufficiently small, and a limsup ε -equilibrium is called a *uniform ε -equilibrium*.

Definition 7 Let $\varepsilon > 0$. A strategy profile σ is a uniform ε -equilibrium if it is a limsup ε -equilibrium and there are $T_0 \in \mathbf{N}$ and $\lambda_0 \in (0, 1)$ such that for every $T \geq T_0$ the strategy profile σ is a T -stage ε -equilibrium and for every $\lambda \in (0, \lambda_0)$ it is a λ -discounted ε -equilibrium.

If for every $\varepsilon > 0$ the game has a T -stage (resp. λ -discounted, limsup, uniform) ε -equilibrium with corresponding payoff g_ε , then any accumulation point of $(g_\varepsilon)_{\varepsilon>0}$ as ε goes to 0 is termed a T -stage (resp. λ -discounted, limsup, uniform) equilibrium payoff.

Zero-Sum Games

A two-player stochastic game is *zero-sum* if $u_1(s, a) + u_2(s, a) = 0$ for every $(s, a) \in SA$. As in matrix games, every two-player zero-sum stochastic game admits at most one equilibrium payoff at every initial state s^1 , which is termed the *value* of the game at the initial state s^1 . Each player's strategy which is part of an ε -equilibrium is termed *ε -optimal*. The definition of ε -equilibrium implies that an ε -optimal strategy guarantees the value up to ε ; for example, in the T -stage evaluation, if σ_1 is an ε -optimal strategy of Player 1, then for every strategy of Player 2 we have

$$\gamma_1^T(s^1, \sigma_1, \sigma_2) \geq v^T(s^1) - \varepsilon,$$

where $v^T(s^1)$ is the T -stage value at s^1 .

In his seminal work, Shapley (1953) presented the model of two-player zero-sum stochastic games with finite state and action spaces and proved the following.

Theorem 8 (Shapley 1953) For every two-player zero-sum stochastic game, the λ -discounted value at every initial state exists. Moreover, both players have λ -discounted 0-optimal stationary strategies.

The λ -discounted value of the game at the initial state s^1 is denoted by $v^\lambda(s^1)$.

Proof Let V be the space of all functions $v: S \rightarrow \mathbf{R}$. For every $v \in V$ and $s \in S$, define a zero-sum matrix game $G_s^\lambda(v)$ as follows:

- The action spaces of the two players are $A_1(s)$ and $A_2(s)$, respectively.
- The payoff function (that Player 2 pays Player 1) is

$$\lambda u_1(s, a) + (1 - \lambda) \sum_{s' \in S} q(s'|s, a) v(s').$$

The game $G_s^\lambda(v)$ captures the situation in which, after the first stage, the game terminates with a terminal payoff $v(s')$, where s' is the state reached after stage 1. Define an operator $\varphi: V \rightarrow V$, termed the *Shapley operator*, as follows:

$$\varphi_s(v) = \text{val}(G_s^\lambda(v)),$$

where $\text{val}(G_s^\lambda(v))$ is the value of the matrix game $G_s^\lambda(v)$. Since the value operator is non-expansive, it follows that the operator φ is contracting: $\|\varphi(v) - \varphi(w)\|_\infty \leq (1 - \lambda)\|v - w\|_\infty$, so that this operator has a unique fixed point $\hat{v}^\lambda$. One can show that the fixed point is the value of the stochastic game, and every strategy σ_i of player i in which he plays, after each finite history $(s^1, a^1, s^2, a^2, \dots, s^t)$, an optimal mixed action in the matrix game $G_{s^t}^\lambda(\hat{v}^\lambda)$ is a λ -discounted 0-optimal strategy in the stochastic game.

Example 9 Consider the two-player zero-sum game with three states: $S = \{s_0, s_1, s_2\}$ that appears in Fig. 1, each entry of the matrix indicates the payoff that Player 2 (the column player) pays Player 1 (the row player; the payoff is in the middle), and the transitions (which are deterministic and are denoted at the top-right corner).

The states s_0 and s_1 are absorbing: once the play reaches one of these states, it never leaves it. State s_2 is nonabsorbing. Stochastic games with a single nonabsorbing state are called absorbing games. For every $v = (v_0, v_1, v_2) \in V = \mathbf{R}^3$, the games $(G_s^\lambda(v))_{s \in S}$ are depicted in Fig. 2.

The unique fixed point $\hat{v}$ of the operator $\text{val}(G^\lambda)$ must satisfy

- $\hat{v}_0 = \text{val}(G_{s_0}^\lambda(\hat{v}))$, so that $v^\lambda(s_0) = \hat{v}_0 = 0$;
- $\hat{v}_1 = \text{val}(G_{s_1}^\lambda(\hat{v}))$, so that $v^\lambda(s_1) = \hat{v}_1 = 1$;
- $\hat{v}_2 = \text{val}(G_{s_2}^\lambda(\hat{v}))$.

Stochastic Games,**Fig. 1** The game in Example 9

	L	R			
T	$0 \quad s^2$	$1 \quad s^1$	T	$1 \quad s^1$	T
B	$1 \quad s^1$	$0 \quad s^0$			
	$State \ s_2$		$State \ s_1$		$State \ s_0$

Stochastic Games,**Fig. 2** The games $(G_s^\lambda(v))_{s \in S}$ in Example 9

	L	R			
T	$(1-\lambda)v_2$	$\lambda+(1-\lambda)v_1$	T	$\lambda+(1-\lambda)v_1$	T
B	$\lambda+(1-\lambda)v_1$	$(1-\lambda)v_0$			
	$The\ game\ G_{s_2}^\lambda$		$The\ game\ G_{s_1}^\lambda$		$The\ game\ G_{s_0}^\lambda$

Stochastic Games,**Fig. 3** The Big Match game in Example 10

	L	R			
T	$0 \quad s^2$	$1 \quad s^2$	T	$1 \quad s^1$	T
B	$1 \quad s^1$	$0 \quad s^0$			
	$State \ s_2$		$State \ s_1$		$State \ s_0$

By Theorem 8 both players have a stationary λ -discounted 0-optimal strategy. Denote by $[x(T), (1 - x)(B)]$ (resp. $[y(L), (1 - y)(R)]$) a mixed action for Player 1 (resp. Player 2) that is part of a λ -discounted 0-optimal strategy at the state s_2 . Since we know that $\hat{v}_0 = 0$ and $\hat{v}_1 = 1$, by playing a completely mixed stationary strategy, Player 1 (resp. Player 2) can guarantee that the λ -discounted payoff is positive (resp. less than 1). Consequently, in his 0-optimal stationary strategy, each player plays a completely mixed action at every stage. This implies that $\hat{v}_2$ is the unique solution of

$$v_2 = y(1 - \lambda)v_2 + (1 - y) = y,$$

so that $v^\lambda(s_2) = \hat{v}_2 = \frac{1 - \sqrt{\lambda}}{1 - \lambda}$. The 0-optimal strategy of Player 2 at state s_2 is $y = \hat{v}_2 = \frac{1 - \sqrt{\lambda}}{1 - \lambda}$, and the 0-optimal strategy of Player 1, $x = \hat{v}_2 = \frac{1 - \sqrt{\lambda}}{1 - \lambda}$, can be found by finding his 0-optimal strategy in $G_{s_2}^\lambda(\hat{v})$.

Bewley and Kohlberg (1976) proved that when the state and action spaces are finite, the function $\lambda \mapsto v^\lambda(s)$, which assigns to every state s and every discount rate λ the λ -discounted value at the initial state s , is a Puiseux function, that is, it has a representation $v^\lambda(s) = \sum_{k=K}^{\infty} a_k \lambda^{k/M}$ that is valid for every $\lambda \in (0, \lambda_0)$ for some $\lambda_0 > 0$, where M is a natural number, K is a nonnegative integer, and $(a_k)_{k=K}^{\infty}$ are real numbers. In particular, the

function $\lambda \mapsto v^\lambda(s)$ is monotone in a neighborhood of 0, and its limit as λ goes to 0 exists. This result turned out to be crucial in subsequent study on games with finitely many states and actions.

Shapley's work has been extended to general state and action spaces; for a recent survey, see Jaśkiewicz and Nowak (2017a). The tools developed in Nowak (2003a), together with a dynamic programming argument, prove that under proper conditions on the payoff function and on the transitions, the two-player zero-sum stochastic game has a T -stage value.

Maitra and Sudderth (1998) proved that the limsup value exists in a very general setup. Their proof follows closely that of Martin (1998) for the determinacy of Blackwell games.

The study of the uniform value emanated from an example, called the “Big Match,” due to Gillette (1957), which was solved by Blackwell and Ferguson (1968).

Example 10 Big Match Consider the stochastic game with two absorbing states and one non-absorbing state that appears in Fig. 3.

Suppose the initial state is s_2 . As long as Player 1 plays T, the play remains at s_2 ; once he plays B, the play moves to either s_0 or s_1 and is effectively terminated. By finding the fixed point of the operator φ , one can show that the λ -discounted value at the

initial state s_2 is $\frac{1}{2}$ and a λ -discounted stationary 0-optimal strategy (i.e., at each stage Player 2 plays L with probability $\frac{1}{2}$ and R with probability $\frac{1}{2}$) for Player 2 is $[\frac{1}{2}(L), \frac{1}{2}(R)]$. Indeed, if Player 1 plays T , then the expected stage payoff is $\frac{1}{2}$ and play remains at s_2 , while if Player 1 plays B , then the game moves to an absorbing state, and the expected stage payoff from that stage onwards is $\frac{1}{2}$. In particular, this strategy guarantees $\frac{1}{2}$ for Player 2 both in the limsup evaluation and uniformly. A λ -discounted 0-optimal strategy for Player 1 is $[\frac{1}{1+\lambda}(T), \frac{1}{1+\lambda}(B)]$.

What can Player 1 guarantee in the limsup evaluation and uniformly? If Player 1 plays the stationary strategy $[x(T), (1-x)(B)]$ that plays at each stage the action T with probability x and the action B with probability $1-x$, then Player 2 has a reply that ensures that the limsup payoff is 0: if $x = 1$ and Player 2 always plays L , the payoff is 0 at each stage; if $x < 1$ and Player 2 always plays R , the payoff is 1 until the play moves to s_0 , and then it is 0 forever. Since Player 1 plays the action B with probability $1-x > 0$ at each stage, the distribution of the stage in which play moves to s_0 is geometric. Therefore, the limsup payoff is 0, and if λ is sufficiently small, the discounted payoff is close to 0.

One can verify that if Player 1 uses a bounded-recall strategy, that is, a strategy that uses only the last k actions that were played, Player 2 has a reply that guarantees that the limsup payoff is 0 and the discounted payoff is close to 0, provided λ is close to 0. Thus, in the limsup payoff, and in the uniform game, finite memory cannot guarantee more than 0 in this game (see also Fortnow and Kimmel (1998)).

To get a limsup payoff higher than 0, Player 1 would like to condition the probability of playing T on the past behavior of Player 2: if in the past Player 2 played the action L more often than the action R , he would have liked to play T with higher probability; if in the past Player 2 played the action R more often than the action L , he would have liked to play B with higher probability. Blackwell and Ferguson (1968) constructed a family of good strategies $\{\sigma_1^M, M \in \mathbb{N}\}$ for Player 1. The parameter M determines the amount that the strategy guarantees: the strategy σ_1^M guarantees a limsup payoff and a discounted payoff of $\frac{M-1}{2M}$, provided the discount

rate is sufficiently low. In other words, Player 1 cannot guarantee $\frac{1}{2}$ but he may guarantee an amount as close to $\frac{1}{2}$ as he wishes by choosing M to be sufficiently large. The strategy σ_1^M is defined as follows: at stage t , play B with probability $\frac{1}{(M+l_t-r_t)^2}$, where l_t is the number of stages up to stage t in which Player 2 played L and r_t is the number of stages up to stage t in which Player 2 played R .

Since $r_t + l_t = t - 1$, one has $r_t - l_t = 2r_t - (t - 1)$. The quantity r_t is the total payoff that Player 1 received in the first $t - 1$ stages if Player 1 played T in those stages (and the game was not absorbed). Thus, this total payoff is a linear function of the difference $r_t - l_t$. When presented this way, the strategy σ_1^M depends on the total payoff up to the current stage. Observe that as r_t increases, $r_t - l_t$ increases as well and the probability to play B decreases.

Mertens and Neyman (1981) generalized the idea presented at the end of Example 10 to stochastic games with finite state and action spaces. (Mertens and Neyman's (1981) result actually holds in every stochastic game that satisfies proper conditions, which are always satisfied when the state and action spaces are finite.)

Theorem 11 *If the state and action spaces of a two-player zero-sum stochastic game are finite, the game has a uniform value $v^0(s)$ at every initial state $s \in S$. Moreover, $v^0(s) = \lim_{\lambda \rightarrow 0} v^\lambda(s) = \lim_{T \rightarrow \infty} v^T(s)$.*


In their proof, Mertens and Neyman describe a uniform ε -optimal strategy. In this strategy the player keeps a parameter, λ_t , which is a fictitious discount rate to use at stage t . This parameter changes at each stage as a function of the stage payoff; if the stage payoff at stage t is high, then $\lambda_{t+1} < \lambda_t$, whereas if the stage payoff at stage t is low, then $\lambda_{t+1} > \lambda_t$. The intuition is as follows. As mentioned before, in stochastic games there are two forces that influence the player's behavior: he tries to get high stage payoffs, while keeping future prospects high (by playing in such a way that the next stage that is reached is favorable). When considering the λ -discounted payoff, there is a clear comparison between the importance of

the two forces: the weight of the stage payoff is λ and the weight of future prospects is $1 - \lambda$; the lower the discount rate, the more weight is given to the future. When considering the limsup value or the uniform value, the weight of the stage payoff is 0. However, if the player never attempts to receive a high stage payoff, the overall payoff in the game will not be high.

To overcome this difficulty, the player keeps a fictitious discount rate; if past payoffs are low and they do not meet the expectation, Player 1 increases the weight of the stage payoff by increasing the fictitious discount rate; if past payoffs are high, Player 1 increases the weight of the future by lowering this fictitious discount rate. This balancing act can be seen in the optimal strategies described in the Big Match: As we calculated above, the λ -discounted 0-optimal strategy of Player 1 in the Big Match is $\left[\frac{1}{1+\lambda}(T), \frac{\lambda}{1+\lambda}(B) \right]$, so that the per-stage probability of the action T increases as λ goes to 0. In the strategy devised by Blackwell and Ferguson (1968), if Player 1 is doing well with the game not having absorbed so far ($r_t - l_t$ is high), he will put higher weight on action T , which corresponds to the optimal λ -discounted stationary strategy with a lower discount rate. On the other hand, if the payoffs have been poor so far ($r_t - l_t$ is negative), he will play action B with higher probability, which corresponds to the optimal λ -discounted stationary strategy with a higher discount rate.

For many years, it had remained an open question as to whether the existence of the uniform value, or even the weaker existence of the *asymptotic value* – i.e., existence of the limit $\lim_{\lambda \rightarrow 0} v_\lambda$ – held if one kept the state space finite but allowed the players to have infinite (but compact) action spaces, under appropriate continuity assumptions on payoffs and transitions. This was finally answered in the negative by Viger (2013) and again by Ziliotto (2016) as a modification of other examples presented by the latter. As elaborated in Sorin and Viger (2015), all these examples can be constructed beginning from simple structures.

Example 12 Consider a two-player zero-sum stochastic game with a single nonabsorbing state s_0 and two absorbing states, s_1 and s_2 , in



Stochastic Games, Fig. 4 The basic game in Example 12

which the action set of each player in state s_0 is $[0, 1]$, the stage payoff in states s_0 and s_1 is α regardless of the actions of the players, and the absorbing payoff in state s_2 is β (see Fig. 4, where $q(s' | s, x)$ represents the multilinear extension of the transition function). If $\alpha < \beta$, Player 1 wants the play to reach state s_2 as soon as possible and Player 2 wants to postpone absorption at s_2 as much as possible, while if $\alpha > \beta$, the goals of the two players are reversed. The λ -discounted payoff under stationary strategy profile $x = (x_1, x_2)$ at the initial state is s_0 .

$$\gamma^\lambda(s_0, x) = \alpha d_\lambda(s_0, x) + \beta(1 - d_\lambda(s_0, x)),$$

where

$$d_\lambda(s_0, x) = \frac{\lambda + (1-\lambda)q(s_2 | s_0, x)}{\lambda + (1-\lambda)(q(s_2 | s_0, x) + q(s_1 | s_0, x))}.$$

Provided the transition function is continuous, the λ -discounted value exists for every discount rate $\lambda > 0$ and


$$v_\lambda(s_0) = \alpha Q_\lambda + \beta(1 - Q_\lambda),$$

where

$$Q_\lambda = \begin{cases} \sup_{\sigma_2} \inf_{\sigma_1} d_\lambda(s_0, \sigma) = \inf_{\sigma_1} \sup_{\sigma_2} d_\lambda(s_0, \sigma) & \text{if } \alpha < \beta, \\ \sup_{\sigma_1} \inf_{\sigma_2} d_\lambda(s_0, \sigma) = \inf_{\sigma_2} \sup_{\sigma_1} d_\lambda(s_0, \sigma) & \text{if } \alpha > \beta, \end{cases}$$

is called the inertia rate. The quantity $d_\lambda(s_0, x)$ is independent of α and β , and the quantity Q_λ depends on α and β only through the identity of the larger number among the two. Consequently, the ε -optimal strategies of the players do not depend on the specific values of α , β , as long as $\alpha < \beta$ (resp. $\alpha > \beta$).

Consider now a game that is composed of two of the games in Fig. 4, where an absorption in one of the subgames leads the play to the other subgame, as appears in Fig. 5. In this game, Player 1 controls the transition in state s_0 , Player 2 controls the transition in state s_2 , and we denote by x_1 and x_2 the stationary strategies of the two players.

Stochastic Games,**Fig. 5** The combined game in Example 12

Denote by Q_λ^1 the inertia rate of the left-hand side part of the game in Fig. 5 and by Q_λ^2 the inertia rate of the right-hand side part of the game in Fig. 5. Using Shapley's result, one can show that the λ -discounted value in states s_0 and s_2 satisfies

$$\begin{aligned} v_\lambda(s_0) &= Q_\lambda^1 \times (-1) + (1 - Q_\lambda^1) \times v_\lambda(s_2), \\ v_\lambda(s_2) &= Q_\lambda^1 \times (+1) + (1 - Q_\lambda^2) \times v_\lambda(s_0), \end{aligned}$$

and therefore

$$v_\lambda(s_0) = \frac{Q_\lambda^2 - Q_\lambda^1 - Q_\lambda^1 Q_\lambda^2}{Q_\lambda^2 + Q_\lambda^1 - Q_\lambda^1 Q_\lambda^2}, \quad v_\lambda(s_2) = \frac{Q_\lambda^2 - Q_\lambda^1 - Q_\lambda^1 Q_\lambda^2}{Q_\lambda^2 + Q_\lambda^1 - Q_\lambda^1 Q_\lambda^2}.$$

In particular, if for each $i = 1, 2$ there is a continuous, positive, and bounded (the notation $g(\lambda) \sim h(\lambda)$ as $\lambda \rightarrow 0$ means $\lim_{\lambda \rightarrow 0} \frac{g(\lambda)}{h(\lambda)} = 1$.) function $f^i: (0, 1] \rightarrow \mathbf{R}$, bounded away from 0 and satisfying $Q_\lambda^i \sim \lambda^r f^i(\lambda)$ as $\lambda \rightarrow 0$, for some $r > 0$, then

$$v_\lambda(s_0) \sim v_\lambda(s_2) \sim \frac{1 - \frac{f^1(\lambda)}{f^2(\lambda)}}{1 + \frac{f^1(\lambda)}{f^2(\lambda)}}.$$

Consequently, to find a stochastic game with no asymptotic value, it becomes a matter to find functions f^1 and f^2 such that Q_λ^1 and Q_λ^2 each has the form above with the same exponent r , one in which $\lim_{\lambda \rightarrow \infty} f^i(\lambda)$ exists for $i = 1$ but not for $i = 2$. One pair of functions (f^1, f^2) that has this property is:

- Assume $q(s_2 | s_0, x) = x$ and $q(s_1 | s_0, x) = x^2$. Minimizing (1), which is now a function of a single variable, gives $x_{1,\lambda} = \sqrt{\frac{\lambda}{1-\lambda}}$, and therefore $Q_\lambda^1 \sim 2\sqrt{\lambda}$.
- Assume $q(s_3 | s_2, y) = y^2$ and $q(s_2 | s_2, y) = y = h(y)$, where h is bounded away from 0, $h'(x) = o(\frac{1}{x})$, and $\lim_{x \rightarrow 0} h(x)$ does not exist,

e.g., $h(x) = 2 + \sin(\ln(-\ln(x)))$. One can compute $x_{2,\lambda} \sim \sqrt{\lambda}$. Since $h'(x) = o(\frac{1}{x})$, we have $h(x_{2,\lambda}) \sim \sqrt{\lambda}$, and substituting this back gives $Q_\lambda^2 \sim \frac{2\sqrt{\lambda}}{h(\lambda)}$.

In light of the example by Vigeral (2013), it is of interest to find classes of stochastic games with infinite action spaces in which the asymptotic value or even the uniform value does exist. A first step in this direction using algebraic tools is presented in Bolte et al. (2015). Further advances in zero-sum stochastic games, and zero-sum dynamic games in general, can be found in Laraki and Sorin (2015).

Multplayer Games

Takahashi (1964) and Fink (1964) extended Shapley's (1953) result to discounted equilibria in nonzero-sum games.

Theorem 13 Every stochastic game with finite state and action spaces has a λ -discounted equilibrium in stationary strategies.

Proof The proof utilizes Kakutani's fixed point theorem (Kakutani, 1941). Let $M = \max_{i,s,a} |u_i(s, a)|$ be a bound on the absolute values of the payoffs. Set $X = \times_{i \in N, s \in S} (\Delta(A_i(s)) \times [-M, M])$. A point $x = (x_{i,s}^A, x_{i,s}^V)_{i \in N, s \in S} \in X$ is a collection of one mixed action and one payoff to each player at every state. For every $v = (v_i)_{i \in N} \in [-M, M]^{N \times S}$ and every $s \in S$, define a matrix game $G_s^\lambda(v)$ as follows:

- The action spaces of each player i is $A_i(s)$.
- The payoff to player i is

$$\lambda u_i(s, a) + (1 - \lambda) \sum_{s' \in S} q(s' | s, a) v_i(s').$$

We define a set-valued function $\varphi: X \rightarrow X$ as follows:

- For every player $i \in N$ and every state $s \in S$, $\varphi_{i,s}^A$ is the set of all best responses of player i to the strategy vector $x_{-i}, s := (x_j, s)_{j \neq i}$ in the game $G_s^\lambda(v)$. That is,

$$\varphi_{i,s}^A(x, v) := \operatorname{argmax}_{y_i, s \in \Delta(A_i(s))} \left(\lambda r_i(s, y_i, s, x_{-i}, s) + (1 - \lambda) \sum_{s' \in S} q(s' | s, y_i, s, x_{-i}, s) v_{i,s'} \right).$$

- For every player $i \in N$ and every state $s \in S$, the quantity $\varphi_{i,s}^V(x_s, v)$ is the payoff for player i in the game $G_s^\lambda(v)$, when the player plays the mixed action profile x_s in the game $G_s^\lambda(v)$. That is,

$$\varphi_s^V(x, v) := \lambda r(s, x_s) + (1 - \lambda) \sum_{s' \in S} q(s' | s, x_s) v_{i,s'}.$$

The set-valued function φ has convex and non-empty values, and its graph is closed, so that by Kakutani's fixed point theorem, it has a fixed point. It turns out that if (x^*, v^*) is a fixed point of φ , then the stationary strategy profile x^* is a stationary λ -discounted equilibrium with corresponding payoff v^* .

This result readily extends to games with discrete countable state spaces, but attempts at proving existence of *stationary* equilibrium when the state space is generally proved elusive. Some deduced the existence of such equilibria in specific classes of games, e.g., Amir (1996) and Horst (2005). Other works establish existence of equilibria in more complex (history-dependent) strategies – see Mertens and Parthasarathy (1987) and Solan (1998) – or in correlated equilibrium, e.g., Nowak and Raghavan (1992); all

these works assume some continuity conditions on the transition function. Recently Levy (2013a) and Levy and McLennan (2015) have demonstrated that general discounted stochastic games may not possess stationary equilibria, presenting examples both in the framework of games with a deterministic transition function, and in the case of absolutely continuous transition function (the latter assumption being undertaken, e.g., in Nowak and Raghavan (1992) and Nowak (2003b)).

As in the case of zero-sum games, a dynamic programming argument coupled with Nash's equilibrium theorem shows that under a strong continuity assumption on the payoff function or on the transition function, a T -stage equilibrium exists.

Little is known regarding the existence of the limsup equilibrium and uniform equilibrium, even when the sets of states and actions are finite. The following classical example, due to Sorin (1986) and coined the “Paris Match,” demonstrates a sort of discontinuity between the sets of discounted and finite-stage equilibria on the one hand and the sets of limsup and uniform equilibria on the other hand.


Example 14 Paris Match Consider the two-player nonzero-sum stochastic game with two absorbing states and one nonabsorbing state, which appears in Fig. 6 and is quite similar to the Big Match.

Suppose the initial state is s_2 . Like in the Big Match, as long as Player 1 plays action T , the play remains at state s_2 ; once he plays action B , the play moves to either state s_0 or s_1 and is effectively terminated. Unlike in zero-sum games, in which the discounted values and finite stage values converge to the uniform value, Sorin (1986) demonstrates that for any discount rate, the only equilibrium payoff is $V := (\frac{1}{2}, \frac{2}{3})$. A similar conclusion holds for the finitely repeated game. However,

Stochastic Games,

Fig. 6 The game in Example 14

	L	R		
T	$(1, 0) s^2$	$(0, 1) s^2$	L	L
B	$(0, 2) s^1$	$(1, 0) s^0$	$(0, 2) s^1$	$(1, 0) s^0$
	State s_2		State s_1	State s_0



Stochastic Games, Fig. 7 Payoffs in the Paris Match

the sets of limsup equilibrium payoffs and uniform equilibrium payoffs coincide with the set of convex combinations of $(\frac{1}{2}, 1)$ and $(\frac{2}{3}, \frac{2}{3})$, which in particular does not include the point V (see Fig. 7).

We sketch the arguments. Fix a discount rate λ . Playing the stationary strategy $[\frac{1}{3}(L), \frac{2}{3}(R)]$ guarantees Player 2 at least $V_2 = \frac{2}{3}$. Since Player 1's payoffs are identical to those in the Big Match, he can guarantee $V_1 = \frac{1}{2}$. Hence, if $W = (W_1, W_2)$ is a λ -discounted equilibrium payoff under equilibrium strategies (σ_1, σ_2) , then it satisfies $W_1 \geq V_1, W_2 \geq V_2$. Standard continuity arguments show that the set of λ -discounted equilibrium payoffs is closed. Assume then that W is a λ -discounted equilibrium payoff with the largest payoff W_2 for Player 2 among all λ -discounted equilibrium payoff. We will show $W_2 \leq V_2$. A similar argument shows that if W_1 is the largest payoff for Player 1 in any equilibrium, then $W_1 \leq V_1$. Together, these implications show that V is the only λ -discounted equilibrium payoff.

It is easy to check that in equilibrium the strategies of the two players play at the first stage a fully mixed action. Let W_L (resp. W_R) denote the expected payoff for Player 2 from the second stage onward (re-normalized) conditional on the action L (resp. R) being played at the first stage. Then W_L

and W_R are themselves λ -discounted equilibrium payoffs; hence both do not exceed W_2 . Denote by p the probability that Player 1 plays the action B at stage 1. As mentioned above, $0 < p < 1$, Player 1 is indifferent among his two actions at the first stage, and consequently the equilibrium condition implies that

$$W_2 = 2p + (1 - p)(1 - \lambda)$$

$$W_L = 0p + (1 - p)(\lambda + (1 - \lambda)W_R).$$

Combining this equation with the condition $W_L, W_R \leq W_2$, a few algebraic manipulations show that $W_2 \leq \frac{2}{3} = V_2$, as required. A similar argument applies to the finitely repeated game.

We turn to the limsup and uniform equilibria. We argue that the set of limsup and uniform equilibrium payoffs coincides with the convex hull of $(\frac{1}{2}, 1)$ and $(\frac{2}{3}, \frac{2}{3})$. Since the payoff function of Player 1 coincides with his payoff function in the Big Match, Player 1's payoff in any uniform equilibrium payoff and in any limsup equilibrium is at least 1. By playing the stationary strategy $[\frac{1}{3}(L), \frac{2}{3}(R)]$, Player 2 guarantees a payoff of $\frac{2}{3}$ in the discounted evaluation and in the limsup evaluation. Consequently, if $W = (W_1, W_2)$ is a limsup equilibrium payoff or a uniform equilibrium payoff, then $W_1 \geq V_1, W_2 \geq V_2$. We next argue that the probability of absorption under any limsup and uniform ε -equilibrium must be close to one. Indeed, to quote Sorin (1986), "If the probability of getting an absorbing payoff on the equilibrium path is less than 1, then after some time Player 1 is essentially playing T ; the corresponding feasible rewards from this stage on are not individually rational [as the payoffs from that stage onward would sum to 1], hence a contradiction."

Conversely, for a payoff $W = (p, 2(1 - p))$ with $\frac{1}{2} \leq p \leq \frac{2}{3}$, let σ_2 be the stationary strategy that plays the mixed action $[(1 - p)(L), p(R)]$ at each stage, and let σ_1^ε be an ε -optimal strategy in the auxiliary zero-sum game that appears in Fig. 8. The auxiliary game is similar to the Big Match (Example 10), and an ε -optimal strategy for Player 1 can be constructed similarly to the construction in that game or using the general construction of Mertens and Neyman (1981). Either of these explicit constructions shows that play under $(\sigma_1^\varepsilon, \sigma_2)$ absorbs with probability 1 and

Stochastic Games,**Fig. 8** The auxiliary game

	L		R			L		L	
T	$p \quad s^2$		$p - 1 s^2$			$-p \quad s^1$		$1 - p s^0$	
B	$-p \quad s^1$		$1 - p s^0$		T				
	State s_2					State s_1		State s_0	

gives the desired payoff. One then verifies using Big Match-type arguments that $(\sigma_1^{\varepsilon}, \sigma_2)$ is a limsup ε -equilibrium and a uniform ε -equilibrium.

The most significant result in the study of the uniform equilibrium so far has been Vieille (2000a,b), who proved that every two-player stochastic game with finite state and action spaces has a uniform ε -equilibrium for every $\varepsilon > 0$. This result has been proved for other classes of stochastic games; see, e.g., Thuijsman and Raghavan (1997), Solan (1999), Solan and Vieille (2001), Simon (2003, 2007, 2012), Altman et al. (2008), Flesch et al. (2007), and Flesch et al. (2008, 2009). Several influential works in this area are Kohlberg (1974), Vrieze and Thuijsman (1989), and Flesch et al. (1997). Most of the papers mentioned above rely on the vanishing discount rate approach, which constructs a uniform ε -equilibrium by studying a sequence of λ -discounted equilibria as the discount rate goes to 0.

For games with general state and action spaces, a limsup equilibrium exists under an ergodicity assumption on the transitions; see, e.g., Nowak (2003b, Remark 4) and Jaśkiewicz and Nowak (2005, 2006).

A particular class of games of interest is those with *perfect information*, i.e., those games in which there are no simultaneous moves, and both players observe past play. Existence of equilibrium in this class was proven by Mertens (1987) in a very general setup of perfect information games. However, it leaves the existence of ε -subgame-perfect equilibrium, that is, a strategy profile inducing ε -equilibria in every subgame. Flesch et al. (2010) show that ε -subgame-perfect equilibrium exists when the payoffs depend in a lower-semicontinuous way on the history of play. (The sequence of plays is endowed with the Tychonoff topology.) Purves and Sudderth (2011) complemented this by proving the

existence of ε -subgame-perfect equilibrium with upper-semicontinuous payoffs. Flesch and Predtetchinski (2017) generalized this result to games in which payoff continuity fails on a sigma-discrete set (a countable union of discrete sets). Flesch et al. (2014) present an example of a perfect information game in which no ε -subgame-perfect equilibrium exists. (Payoffs are Borel and described explicitly.) For a recent survey on nonzero-sum stochastic games, the reader is referred to Jaśkiewicz and Nowak (2017b).

Correlated Equilibrium

The notion of correlated equilibrium was introduced by Aumann (1974, 1987); see also Forges (2007). A *correlated equilibrium* of a static game is an equilibrium of an extended game, in which each player receives at the outset of the game a private signal such that the vector of signals is chosen according to a known joint probability distribution. In repeated interactions, such as in stochastic games, there are three natural notions of correlated equilibria: (a) each player receives one signal at the outset of the game (*normal-form correlated equilibrium*); (b) each player receives a signal at each stage (*extensive-form correlated equilibrium*); and (c) at every stage the players observe a public signal that is, without loss of generality, taken to be uniformly distributed in $[0, 1]$ (*extensive-form public correlated equilibrium*). It follows from Forges (1990) that when the state and action sets are finite, the set of all correlated T -stage equilibrium payoffs (either normal form or extensive form) is a polytope.

Nowak and Raghavan (1992) proved the existence of stationary discounted extensive-form public correlated equilibrium under weak conditions on the state and action spaces. Roughly, their approach is to apply Kakutani's fixed point theorem to the set-valued function that assigns to each

game $G_s^\lambda(v)$ the set of all correlated equilibrium payoffs in this game, which is convex and compact. This approach has since been generalized by Duggan (2012) to equilibrium in games in which the state component contains enough “noise,” which is used to replace the public signal.

Solan and Vieille (2002) proved the existence of an extensive-form correlated uniform equilibrium payoff when the state and action spaces are finite. Their approach is to let each player play his uniform optimal strategy in a zero-sum game in which all other players try to minimize his payoff. Existence of a normal-form correlated equilibrium was proved for the class of absorbing games (Solan and Vohra, 2002).

Solan (2001) characterized the set of extensive-form correlated equilibrium payoffs for general state and action spaces and a general evaluation on the stage payoffs and provided a sufficient condition that ensures that the set of normal-form correlated equilibrium payoffs coincides with the set of extensive-form correlated equilibrium payoffs. Using these techniques, Mashiach-Yaakovi (2015) showed the existence of extensive-form correlated equilibrium in stochastic games when the payoffs are given by general bounded Borel functions.

Imperfect Monitoring

So far it has been assumed that at each stage the players know the past play. There are cases in which this assumption is too strong; in some cases players do not know the complete description of the current state (Examples 3 and 4), and in others players do not fully observe the actions of all other players (Examples 2, 3, and 4). For a most general description of stochastic games, see Mertens et al. (2015, Chapter IV) and Coulomb (2003b). In this model, at every stage each player observes a private signal, which depends on the current state and on the actions taken by the players at the current stage.

The following observation can be used to show the existence of equilibrium in some classes of games, in particular for discounted equilibrium,

the T -stage equilibrium, or the limsup equilibrium: An alternative description of the game is to view the strategies as being chosen at the onset of the game, hence reducing the stochastic game to a one-shot game in which the actions are the strategies in the stochastic games. This space has a natural topological structure in which two strategies are close if they approximately agree for a long initial period. This reduction is feasible regardless of the information structure. The discounted and T -stage payoffs are continuous with respect to this structure, and hence Nash’s equilibrium theorem applies and allows us to deduce the existence of equilibrium (see, e.g., Altman and Solan, 2009). This approach may be also used for the limsup equilibrium under a proper ergodicity condition.

Whenever there exists an equilibrium in stationary strategies (e.g., a discounted equilibrium in games with finitely many states and actions), the only information that players need in order to follow the equilibrium strategies is the current state. In particular, they need not observe past actions of the other players. As we now show, in the Big Match game the limsup value and the uniform value may fail to exist when each player does not observe the past actions of the other player.

Example 10 Continued Assume that no player observes the actions of the other player, and assume that the initial state is s_2 . Player 2 can still guarantee $\frac{1}{2}$ in the limsup evaluation by playing the stationary strategy $[\frac{1}{2}(L), \frac{1}{2}(R)]$. One can show that for every strategy of Player 2, Player 1 has a reply such that the limsup payoff is at least $\frac{1}{2}$. In other words, $\inf_{\sigma_2} \sup_{\sigma_1} \gamma^\infty(s_2, \sigma_1, \sigma_2) = \frac{1}{2}$.

We now argue that $\sup_{\sigma_1} \inf_{\sigma_2} \gamma^\infty(s_2, \sigma_1, \sigma_2) = 0$. Indeed, fix a strategy σ_1 for Player 1 and $\varepsilon > 0$. Let $\theta \in \mathbf{N}$ be sufficiently large such that the probability that under σ_1 Player 1 plays action B for the first time after stage θ is at most ε . If no such θ exists, then absorption occurs a.s., so the best response of Player 2 is to always play action R . Observe that as t increases, the probability that Player 1 plays action B for the first time after stage t decreases to 0, so that such θ exists. Consider the following strategy σ_2 of Player 2:

play action R up to stage θ and play L from stage $t + 1$ and on. By the definition of θ , Player 1 plays action B either before or at stage θ , and then the game moves to state s_0 , and the payoff is 0 at each stage thereafter; or Player 1 plays action T at each stage, and then the stage payoff after stage θ is 0, or, with probability less than ε , Player 1 plays action B for the first time after stage θ , the play moves to state s_1 , and the payoff is 1 thereafter. Thus, the limsup payoff is at most ε .

A similar analysis shows that for every $0 < \lambda < 1$,

$$\inf_{\sigma_2} \sup_{\sigma_1} \gamma^\lambda(s_2, \sigma_1, \sigma_2) = \frac{1}{2},$$

however

$$\sup_{\sigma_2} \inf_{\sigma_1} \lim_{\lambda \rightarrow 0} \gamma^\lambda(s_2, \sigma_1, \sigma_2) = 0,$$

so that the uniform value does not exist as well.

This example shows that in general the limsup value and the uniform value need not exist when the players do not observe past play. Though in general the value (and therefore also an equilibrium) need not exist, in many classes of stochastic games, the value or an equilibrium does exist, even in the presence of imperfect monitoring.

Rosenberg et al. (2002) and Renault (2011) showed that the uniform value exists in the one player setup (Markov decision problem), in which the player receives partial information regarding the current state. Thus, a single decision-maker who faces a dynamic situation and does not fully observe the state of the environment can play in such a way that guarantees high payoff, provided the interaction is sufficiently long or the discount rate is sufficiently low.

Altman et al. (2005, 2008) studied stochastic games in which each player has a “private” state, which only he can observe, and the state of the world is composed of the vector of private states; each player also does not observe the action of the other players. Such games arise naturally in wireless communication (see Altman et al., 2005); take, for example, several mobiles which periodically send information to a base station. The private state of a mobile may depend, e.g., on its exact physical environment, and it determines the power attenuation between the mobile and the base station. The throughput (the amount of bits per second) that a mobile can send to the base

station depends on the power attenuations of all the mobiles. Finally, the stage payoff is the stage power consumption.

Rosenberg et al. (2009) studied the extreme case of two-player zero-sum games in which the players observe neither the current state nor the action of the other player and proved that the uniform value does exist in two classes of games, which capture the behavior of certain communication protocols. Classes of games in which the actions are observed but the state is not observed were studied, e.g., by Sorin (1984, 1985), Sorin and Zamir (1991), Krausz and Rieder (1997), Flesch et al. (2003), Neyman (2008), Rosenberg et al. (2004), Renault (2006, 2012), and Gensbittel and Renault (2015). For additional results, see Rosenberg et al. (2003), Coulomb (2003a,c), and Sorin (2003).

Until recently, several conjectures raised by Mertens (1987) had remained open in the general model of zero-sum repeated games that allows for imperfect observation of states and actions. Primarily, it was conjectured that when the number of states, actions, and signals is finite, the limits $\lim_{\lambda \rightarrow 0} v_\lambda$ and $\lim_{n \rightarrow \infty} v_n$ exist and are equal. This conjecture was recently shown to be false by Ziliotto (2016). Ziliotto (2016) presents an example of a game with symmetric information (i.e., actions are observed, and consequently at every stage, both players have the same beliefs over the set of states) in which $\lim_{\lambda \rightarrow 0} v_\lambda$ does not exist. In fact, in that example, in each state only one of the players controls the payoffs and transitions, and although states are not observed directly, payoffs are known. Ziliotto (2016) then shows how the example can be modified so that neither $\lim_{\lambda \rightarrow 0} v_\lambda$ nor $\lim_{n \rightarrow \infty} v_n$ exist. Lastly, Ziliotto (2016) presents examples of a state-blind stochastic game (players observe actions but get no information about the state) and an example of a game with one state-blind player (i.e., one player knows the states, but the other receives no information about it) in which the asymptotic value does not exist. See also Sorin and Vigeral (2015).

Another type of incomplete information is the presence of information lag, that is, a delay in learning about the actions of an opponent or of the state evolution; see Levy (2012) and the

references therein. Yet another interesting direction that introduces uncertainty of a different kind is to observe games of finite length but with uncertainty present as to how long the game will proceed. Neyman and Sorin (2010) studies such game where the probabilistic information about the duration is public (i.e., known to both players), deduces a recursive formula for the value of such a game analogous to that of the value for a game of fixed length, and establishes convergence of the value as the expectation of the duration goes to infinity. See reference therein for previous works on asymmetric-information uncertain duration processes in repeated prisoners' dilemma and repeated games of incomplete information.

Folk Theorems

Another direction of study is the topic of *Folk Theorems*. These results, dating back to Aumann and Shapley (1976) and Rubinstein (1979), originated in the study of ordinary repeated games (i.e., a single state played repeatedly) and attempt to characterize the set of possible equilibrium payoffs as the players become more and more patient (i.e., as the discount rate λ goes to 0). If the payoff in the repeated game is given by a function $u(\cdot)$, assigning to each action profile a in the set of action profiles $A = \prod_{i \in N} A_i$ a payoff vector $u(a)$, then the set of *feasible* payoff vectors is given by

$$F = co\{u(a) \mid a \in A\},$$

where $co(\cdot)$ denotes the convex hull. A payoff vector $v = (v_1, \dots, v_n)$ is said to be (strictly) *individually rational* if for every player i ,

$$v_i > m_i := \min_{x_{-i} \in \prod_{j \neq i} \Delta(A_j)} \max_{a_i \in A_i} u_i(a_i, x_{-i});$$

i.e., if it is higher than the minmax value in mixed strategies for all players. The folk theorem for repeated games (with perfect monitoring) states that for every feasible and individually rational payoff vector v , there is a discount rate λ_0 such that if $\lambda \leq \lambda_0$, v is an equilibrium payoff of the λ -discounted game. In fact, assuming the set of feasible and individually rational payoff vectors

has nonempty interior (ruling out, e.g., the case when two agents have identical payoffs), Fudenberg and Maskin (1986) show that all such payoffs can arise in subgame-perfect equilibria as well for small enough discount rate. This result was extended to the case of public imperfect monitoring of actions (i.e., players do not observe others' actions, but rather the action profile determines a distribution over public signals), assuming that the public signals about the action profiles satisfy conditions that effectively allow, in the long-run, detection of deviations and the identity of the deviator, known as "full rank conditions"; see Fudenberg et al. (1994).

Dutta (1995) extended the classical folk theorem to stochastic games, both for the discounted evaluation and for the T -stage evaluation. The λ -discounted feasible set of payoffs, which depends on the initial state s , includes the payoffs that can be achieved in the entire play of the game and not just in the one-shot game:

$$F^\lambda(s) := co(\{\gamma^\lambda(s, \sigma) \mid \sigma \text{ is any strategy profile}\}),$$

where $\gamma^\lambda(s, \sigma) := (\gamma_i^\lambda(s, \sigma))_{i \in N}$ is the vector of payoffs. It is shown there that it suffices to take the convex hull of payoffs of pure stationary profile payoffs. Dutta (1995) similarly defines the feasible set for the long-run average stage game as the convex hull of limits of payoffs in the T -stage games as T goes to infinity:

$$F(s) := co(\{\lim_{T \rightarrow \infty} \gamma^T(s, \sigma) \mid \sigma \text{ is any strategy profile}\}),$$

where $\lim_{T \rightarrow \infty} \gamma^T(s, \sigma)$ refers to the set of accumulation points. In particular, a single strategy profile may yield multiple points in $F(s)$. Individual rationality as well refers to being higher than the minmax of the entire game; i.e., the vector $v \in \mathbf{R}^N$ is (strictly) individually rational in the λ -discounted game if

$$v^i > m^i(s, \lambda) := \inf_{\sigma_{-i}} \sup_{\sigma_i} \gamma_i^\lambda(s, \sigma)$$

and is (strictly) individually rational in the long-run average game if

$$v^i > m^i(s) := \inf_{\sigma_{-i}} \sup_{\sigma_i} \liminf_{T \rightarrow \infty} \gamma_i^T(s, \sigma).$$

To derive folk theorems, Dutta (1995) assumes *asymptotic state independence*; that is, the limits $\lim_{\lambda \rightarrow 0} m^i(s, \lambda)$ for $i \in N$ and $\lim_{\lambda \rightarrow 0} F^\lambda(s)$ are independent of the state. This assumption holds,

for example, in the case of *irreducible stochastic games* with finite state space, i.e., games in which any state will be reached eventually starting at any other state, regardless of the strategies used by the players. Assuming the full dimensionality of the set of feasible payoffs, any long-run average feasible and individually rational payoff vector corresponds to a subgame-perfect equilibrium when the players are sufficiently patient.

More recently, there have been attempts to generalize the imperfect public monitoring framework of Fudenberg et al. (1994) to stochastic games. This has been achieved by Fudenberg and Yamamoto (2011) and Hörner et al. (2011), under the assumptions of asymptotic state independence, full dimensionality of payoffs, and the “full rank conditions” of the public signals; the latter work also discusses algorithms for computing such strategies. Peski and Wiseman (2015) derive similar results, but rather than assuming players become more and more patient, they assume that the duration of each stage of play becomes shorter. Aiba (2014) establishes a folk theorem for irreducible stochastic games with private almost-perfect monitoring of actions, extending the results Hörner and Olszewski (2006) have established for repeated games.

Algorithms

There has been extensive work on algorithms for computing the value and optimal strategies (and, in some cases, equilibria) of stochastic games, particularly in light of their important applications and the increased interest of computer scientists in game-theoretic questions (see Nisan et al. (2007) or Papadimitriou (2007)).

It is well known that the value of a two-player zero-sum matrix game and optimal strategies for the two players can be calculated efficiently using a linear program. Equilibria in two-player nonzero-sum games can be calculated by the Lemke-Howson algorithm, which is usually efficient; however, its worst running time is exponential in the number of pure strategies of the players (Savani and von Stengel 2004). Unfortunately, to date there are no efficient algorithms to calculate either the

value in zero-sum stochastic game or equilibria in nonzero-sum games. Moreover, in Example 9 the discounted value may be irrational for rational discount rates, even though the data of the game (payoffs and transitions) are rational, so it is not clear whether linear programming methods can be used to calculate the value of a stochastic game. Nevertheless, linear programming methods were used to calculate the discounted and uniform value of several classes of stochastic games; see Filar and Vrieze (1996) and Raghavan and Syed (2002). Other methods that were used to calculate the value or equilibria in discounted stochastic games include fictitious play (Vrieze and Tijds 1982), value iterates, policy improvement, and general methods to find the maximum of a function (see Filar and Vrieze (1996) and Raghavan and Syed (2003)), a homotopy method (Herings and Peeters, 2004), and algorithms to solve sentences in formal logic (Chatterjee et al. (2008) and Solan and Vieille (2010)).

There are two related questions one can ask from a computational point of view: The complexity of finding the value and the complexity of finding ε -optimal strategies. The earliest study in this venue is Condon (1992), who concentrates on a subclass of zero-sum games termed *simple stochastic games*. In such games, there are absorbing vertices (*sinks*) with payoffs 0 or 1, and at all other vertices, reward is 0, and either Player 1, or Player 2, or Nature chooses which of possible neighbors to continue on to. Simple stochastic games are games with perfect information; hence both players have pure and stationary limsup optimal strategies. Condon (1992) studies the complexity of finding whether the value of the game is larger than some constant a and shows that this problem lies in the intersection of the complexity class NP , those solvable in polynomial time using randomized algorithms, and $coNP$, those whose negation lie in NP . It is also shown that if transitions are controlled by only one player and Nature, or by the players but not by Nature, then the problem is solvable in polynomial time. Andersson and Miltersen (2009) and the references therein show that solving simple stochastic games (either for the value or for ε -optimal strategies) is as computationally complex (i.e.,

polynomial-time reducible to/form) solving either the discounted or limiting average games with perfect information in general.

Chatterjee et al. (2008) show that the uniform value of a finite-state two-player zero-sum stochastic game with limit-average payoff can be approximated to within ε in time exponential in a polynomial in the size of the game times a polynomial in the logarithm of $1/\varepsilon$. Hansen et al. (2011) give algorithms for solving stochastic games with either discounted or limsup evaluations. The run-times are in general quite long; however, when the number of states is fixed, the algorithms run in polynomial time. Hansen et al. (2016) give algorithms for computing ε -optimal strategies which use $O(\log \log T)$ space, T being the stage, improving on the previous $O(\log T)$ space required by the strategies demonstrated earlier in this paper. In a somewhat different strain, Bertrand et al. (2009) study stochastic games with imperfect monitoring of actions and public signals. Instead of studying the discounted or limsup evaluations, they set winning criteria such as *reachability/safety* (i.e., reaching a specific subset of states eventually/never) or *Büchi/co-Büchi* (i.e., reaching a specific subset of states infinitely many times/at most finitely many times). Giving such an objective for, say, Player 1, the problem is to determine whether Player 1 has a strategy that wins almost surely. In general, they show that solving reachability games is $2EXP$ $TIME$ -complete ($2EXP$ $TIME$ is the class of problems solvable in time $O(2^{2^{p(n)}})$ for some polynomial $p(\cdot)$), as is solving stochastic games with the Büchi winning criterion. Surprisingly, games with the co-Büchi winning criterion are in general *undecidable*, i.e., cannot be solved by a Turing machine.

Continuous-Time Games

The model we have presented is played in *discrete time*. A natural variation is to model the game in *continuous time*. In such a model, the states and the actions of the players may change at any time. As argued by, e.g., Neyman (2012), this could be the natural framework for modeling a wide range

of phenomena, from the study of occurrence of financial crises, in which the state of the economy can change drastically in a very short time, to sports matches like soccer, in which the score can change in a split second.

The earliest model of continuous-time stochastic games (called there *Markov games*) goes back to Zachrisson (1964), who studied zero-sum games played on a fixed interval of time $[0, T]$. In such games, the payoffs are given by integration over the play of the game of the instantaneous payoffs: i.e., given the running payoff which assigns to every state s , every action profile a , and every player i an instantaneous payoff $u_i(s, a)$, the total payoff is given by

$$\int_0^T u_i(s^t, a^t) dt,$$

where s^t and a^t are the state and action profile at time t . The current state and action profile determine the *rate* of transition. That is, the transitions occur in “jumps,” and when the state at time t is s^t and action profile a^t is played from time t to $t + \delta$, the probability of moving to state s' between time t and $t + \delta$ is approximately $\delta q(s' \mid s^t, a^t)$. Zachrisson (1964) concentrates on *Markovian* strategies, which are a function of only state and time, and shows in particular that the value of the zero-sum games exists, as well as do optimal Markovian strategies, when the action spaces are compact and convex and the transitions and payoffs are multilinear in the actions of the players. This assumption holds in particular when the action a_i^t of player i at time t is in fact a mixed action over a finite set of actions.

Until recently, this model has received very little attention, until being revived by Neyman (2012), who considers multiplayer stochastic games with infinite-time horizon and continuous time. As had been well known from other models of continuous-time games (e.g., differential games; see Friedman (1971)), there is some difficulty in defining strategies in continuous time. Allowing a player's strategy at time t to depend on everything that has transpired up to precisely time t is known to not induce a well-defined probability distribution over possible plays of the game. A technique commonly used to get around this problem in the differential game

literature is to allow actions chosen at time t to depend on what had occurred up to time $t - \delta$. One possible such concept is *nonanticipating strategies with delay*, which dictates that a strategy of a player is accompanied by an increasing sequence of stopping time $(\tau_k)_{k \in \mathbb{N}}$, and the behavior of the player between times τ_k and τ_{k+1} depends only on the play before time τ_k .

Under this framework, Neyman (2012) establishes for zero-sum games the existence of the discounted value with stationary optimal strategies and the existence of the uniform value via techniques quite similar to those used in discrete-time games in Mertens and Neyman (1981). Turning to multiplayer games, he uses these and tools used in establishing existence of correlated equilibria in discrete-time games by Solan and Vieille (2002) to deduce the existence of uniform equilibrium in the continuous-time game. An accompanying paper Neyman (2013) shows in what sense, the continuous-time game can be approximated by discrete-time games.

A few other works on continuous-time stochastic games have appeared. Guo and Hernandez-Lerma (2005) establish the existence of discounted equilibrium under some assumptions. Levy (2013b), which, like Zachrisson (1964), studies finite horizon games with Markovian strategies, establishes the existence of correlated equilibrium while characterizing both equilibria and correlated equilibria via Hamilton-Jacobi-Bellman differential inclusions, as well as drawing links between the continuous-time games and discrete-time approximations. Lovo and Tomala (2015) study *stochastic revision games* in which the states and actions of the players can be changed at random times that follow a Poisson process.

In addition, there have been stochastic generalizations of classic differential games. In differential games, the state is a continuous parameter, not discrete, and evolves constantly, with its rate of change being a function of current state and actions:

$$dx = f(x, u)dt$$

where u is the strategy profile and x is the state. Stochastic differential games generalize this by adding a stochastic component to the state evolution via the addition of a noise induced by a standard Brownian motion $(W_t)_{t \geq 0}$ and like differential games are typically studied on a finite horizon:

$$dx = f(x, u)dt + \sigma(x, u)dW_t.$$

Fleming and Souganidis (1989) show the existence of the value for such games under fairly general conditions although, as is standard, assume the *Isaac's condition*, essentially requiring that the minmax and maxmin of an appropriate expression weighing the running payoffs and transitions are equal. They also characterize the value function as a *viscosity solution* of an appropriate differential equation; these types of solutions are generalizations of classical solutions to differential equations and are now ubiquitous in the study of differential games and stochastic optimal control. See Ramachandran and Tsokos (2012) for further results on stochastic differential games.

Finally, we mention the model of continuous time *Dynkin* games, introduced by Dynkin (1967). (Discrete-time Dynkin games are similarly defined and serve a natural generalization of *stopping problems*; see, e.g., Neveu (1975). For some results in this topic, see, e.g., Yasuda (1985), Morimoto (1986), Mamer (1987), Ohtsubo (1987, 1991), Rosenberg et al. (2001), and Solan and Shmaya (2004).) In such games, each player may choose at any time to stop the game. The decisions may be conditioned on information acquired as time goes by, via an increasing right-continuous filtration $(\mathcal{F}_t)_{t \geq 0}$ on a probability space $(\Omega, \mathcal{A}, P)$; hence a (pure) strategy is a $(\mathcal{F}_t)_{t \geq 0}$ -adapted stopping time. In the two-player case (the general case can analogously be defined), let $(X_i, Y_i, Z_i)_{i=1,2}$ be $(\mathcal{F}_t)_{t \geq 0}$ -adapted processes and $(\xi_i)_{i=1,2}$ be bounded real-valued $(V_t \geq 0, \mathcal{F}_t)$ -measurable functions. The payoffs for player i when pure strategies τ_1, τ_2 are used are

$$\gamma_i(\tau_1, \tau_2) = E[X_i(\tau_1)1_{\tau_1 < \tau_2} + Y_i(\tau_2)1_{\tau_1 > \tau_2} + Z_i(\tau_1)1_{\tau_1 = \tau_2} + \xi_i 1_{\tau_1 = \tau_2 = \infty}].$$

Randomized stopping times are then measurable mappings $\varphi_i: [0, 1] \times \Omega \rightarrow [0, \infty]$ such that for each $r \in [0, 1]$, $\varphi_i(r, \cdot): \Omega \rightarrow [0, \infty]$ is a stopping time and the associated payoff under a profile of mixed strategies is

$$\gamma_i(\phi_1, \phi_2) = \int_{[0, 1]^2} \gamma_i(\phi_1(r, \cdot), \phi_2(s, \cdot)) dr ds.$$

Laraki and Solan (2005) demonstrated that two-player zero-sum Dynkin games have a value, although may only have ε -optimal strategies for $\varepsilon > 0$. Laraki and Solan (2013) generalized this to show that nonzero-sum two-player Dynkin

game possesses ε -equilibria. Laraki et al. (2005) showed that three-player games need not possess ε -equilibria, even when the processes $(X_i, Y_i, Z_i)_{i=1,2}$ and the functions $(\xi_i)_{i=1,2}$ are all constant. See Kifer (2013) for a survey of Dynkin games and applications in finance.

Additional and Future Directions

The research on stochastic games extends to additional directions than those mentioned in earlier sections. Approximation of games with infinite state and action spaces by finite games was discussed by Whitt (1980) and further developed by Nowak (1985). For hybrid models that include both discrete-time and continuous-time aspects, see, e.g., Başar and Olsder (1995) and Altman and Gaitsgory (1995).

Among the many directions of future research in this area, we will mention here but a few. One challenging question is the existence of a uniform equilibrium and a limsup equilibrium in multi-player stochastic games with finite state and action spaces. Another active area is to establish the existence of the uniform value or uniform maxmin in classes of stochastic games with imperfect monitoring. A third direction concerns the identification of applications that can be recast in the framework of stochastic games and that can be successfully analyzed using the theoretical tools that the literature developed. Another problem that is of interest is the characterization of approachable and excludable sets in stochastic games with vector payoffs (see Blackwell (1956) for the presentation of matrix games with vector payoffs and Milman (2006) and Flesch et al. (2016) for partial results regarding this problem).

Bibliography

Primary Literature

- Aiba K (2014) A folk theorem for stochastic games with private almost-perfect monitoring. *Games Econ Behav* 86:58–66
- Altman E (2005) Applications of dynamic games in queues. *Adv Dyn Games* 7:309–342. Birkhauser
- Altman E, Gaitsgory VA (1995) A hybrid (differential-stochastic) zero-sum game with fast stochastic part. *Ann Int Soc Dyn Games* 3:47–59
- Altman E, Solan E (2009) Constrained games: the impact of the attitude to Adversary's constraints. *IEEE Trans Autom Control* 54:2435–2440
- Altman E, Avrachenkov K, Marquez R, Miller G (2005) Zero-sum Constrained stochastic games with independent state processes. *Math Meth Oper Res* 62:375–386
- Altman E, Avrachenkov K, Bonneau N, Debbah M, El-Azouzi R, Sadoc MD (2008) Constrained cost-coupled stochastic games with independent state processes. *Oper Res Lett* 36:160–164
- Amir R (1996) Continuous stochastic games of capital accumulation with convex transitions. *Games Econ Behav* 15:111–131
- Andersson D, Miltersen PB (2009) The complexity of solving stochastic games on graphs. In: Dong Y, Du DZ, Ibarra O (eds) *Algorithms and computation. ISAAC 2009, Lecture notes in computer science*, vol 5878. Springer, Berlin/Heidelberg
- Aumann RJ (1974) Subjectivity and correlation in randomized strategies. *J Math Econ* 1:67–96
- Aumann RJ (1987) Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55:1–18
- Aumann RJ, Shapley L (1976) Long term competition: a game theoretical analysis. Mimeo, Hebrew University
- Bertrand N, Genest B, Gimbert H (2009) Qualitative determinacy and decidability of stochastic games with signals. In: 24th annual IEEE symposium on logic in computer science (LICS'09), pp 319–328
- Bewley T, Kohlberg E (1976) The asymptotic theory of stochastic games. *Math Oper Res* 1:197–208
- Blackwell D (1956) An analog of the minimax theorem for vector payoffs. *Pacific J Math* 6:1–8
- Blackwell D, Ferguson TS (1968) The big match. *Ann Math Stat* 39:159–163
- Bolte J, Gaubert S, Vigeral G (2015) Definable zero-sum stochastic games. *Math Oper Res* 40:171–191
- Chari V, Kehoe P (1990) Sustainable plans. *J Polit Econ* 98:783–802
- Chatterjee K, Majumdar R, Henzinger TA (2008) Stochastic limit-average games are in EXPTIME. *Int J Game Theory* 37:219–234
- Condon A (1992) The complexity of stochastic games. *Inf Comput* 96:203–224
- Coulomb JM (2003a) Absorbing games with a signalling structure. In: Neyman A, Sorin S (eds) *Stochastic games and applications*, NATO science series. Kluwer, Dordrecht, pp 335–355
- Coulomb JM (2003b) Games with a recursive structure. In: Neyman A, Sorin S (eds) *Stochastic games and applications*, NATO science series. Kluwer, Dordrecht, pp 427–442
- Coulomb JM (2003c) Stochastic games without perfect monitoring. *Int J Game Theory* 32:73–96
- Duggan J (2012) Noisy stochastic games. *Econometrica* 80:2017–2045
- Dutta P (1995) A folk theorem for stochastic games. *J Econ Theory* 66:1–32
- Dutta P, Sundaram RK (1992) Markovian equilibrium in a class of stochastic games: existence theorems for discounted and undiscounted models. *Economic Theory* 2:197–214

- Dutta P, Sundaram RK (1993) The tragedy of the commons? *Economic Theory* 3:413–426
- Dynkin EB (1967) Game variant of a problem on optimal stopping. *Soviet Math Dokl* 10:270–274
- Filar JA, Vrieze K (1996) *Competitive markov decision processes*. Springer-Verlag, New York
- Fink AM (1964) Equilibrium in a stochastic n -person game. *J Sci Hiroshima Univ* 28:89–93
- Fleming WH, Souganidis PE (1989) On the existence of value functions of two-player zero-sum stochastic differential games. *Indiana Univ Math J* 38:293–314
- Flesch J, Predtetchinski A (2017, forthcoming) A characterization of subgame-perfect equilibrium plays in borel games of perfect information. *Math Oper Res*
- Flesch J, Thuijsman F, Vrieze K (1997) Cyclic Markov equilibria in stochastic games. *Int J Game Theory* 26:303–314
- Flesch J, Thuijsman F, Vrieze OJ (2003) Stochastic games with non-observable actions. *Math Meth Oper Res* 58:459–475
- Flesch J, Thuijsman F, Vrieze OJ (2007) Stochastic games with additive transitions. *Eur J Oper Res* 179:483–497
- Flesch J, Schoenmakers G, Vrieze K (2008) Stochastic games on a product state space. *Math Oper Res* 33:403–420
- Flesch J, Schoenmakers G, Vrieze K (2009) Stochastic games on a product state space: the periodic case. *Int J Game Theory* 38:263–289
- Flesch J, Kuipers J, Mashiah-Yaakovi A, Schoenmakers G, Solan E, Vrieze K (2010) Perfect-information games with lower-semi-continuous payoffs. *Math Oper Res* 35:742–755
- Flesch J, Kuipers J, Mashiah-Yaakovi A, Schoenmakers G, Shmaya E, Solan E, Vrieze K (2014) Non-existence of subgame-perfect ε -equilibrium in perfect information games with infinite horizon. *Int J Game Theory* 43:945–951
- Flesch J, Laraki R, Perchet V (2016) Online learning and Blackwell approachability in quitting games. *J Mach Learn Res* 49:941–942. *Proceedings of the 29th conference on learning theory*
- Forges F (1990) Universal mechanisms. *Econometrica* 58:1341–1364
- Fortnow L, Kimmel P (1998) Beating a finite automaton in the big match. In: *Proceedings of the 7th conference on theoretical aspects of rationality and knowledge*. Morgan Kaufmann, San Francisco, pp 225–234
- Fudenberg D, Maskin E (1986) The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* 54:533–554
- Fudenberg D, Yamamoto Y (2011) The folk theorem for irreducible stochastic games with imperfect public monitoring. *J Econ Theory* 146:1664–1683
- Fudenberg D, Levine D, Maskin E (1994) The folk theorem with imperfect public information. *Econometrica* 62:997–1039
- Gensbittel F, Renault J (2015) The value of Markov chain games with Lack of information on both sides. *Math Oper Res* 40:820–841
- Gillette D (1957) *Stochastic games with zero stop probabilities*, Contributions to the theory of games, vol 3. Princeton University Press, Princeton
- Guo X, Hernandez-Lerma O (2005) Nonzero-sum games for continuous-time Markov chains with unbounded discounted payoffs. *J Appl Probab* 42:303320
- Hansen KA, Kouck M, Lauritzen N, Miltersen PB, Tsigaridas E (2011) Exact algorithms for solving stochastic games. In: *Proceedings of the 43rd annual ACM symposium on theory of computing (STOC'11)*. Association for Computing Machinery, New York, pp 205–214
- Hansen KA, Ibsen-Jenses R, Kouck M (2016) *Algorithmic Game Theory - 9th International Symposium, SAGT 2016*, Proceedings, held in Liverpool, UK; Springer-Verlag, Berlin Heidelberg
- Herings JJP, Peeters RJAP (2004) Stationary equilibria in stochastic games: structure, selection, and computation. *J Econ Theory* 118:32–60
- Hörner J, Olszewski W (2006) The folk theorem for games with private almost-perfect monitoring. *Econometrica* 74:1499–1544
- Hörner J, Sugaya T, Takahashi S, Vieille N (2011) Recursive methods in discounted stochastic games: an algorithm for $\delta \rightarrow 1$; and a folk theorem. *Econometrica* 79:1277–1318
- Horst U (2005) Stationary equilibria in discounted stochastic games with weakly interacting players. *Games Econ Behav* 51:83–108
- Jaśkiewicz A, Nowak AS (2005) Nonzero-sum semi-Markov games with the expected average payoffs. *Math Meth Oper Res* 62:23–40
- Jaśkiewicz A, Nowak AS (2006) Zero-sum ergodic stochastic games with feller transition probabilities. *SIAM J Control Optim* 45:773–789
- Kakutani S (1941) A generalization of Brouwers fixed point theorem. *Duke Math J* 8:457–459
- Kohlberg E (1974) Repeated games with absorbing states. *Ann Stat* 2:724–738
- Krausz A, Rieder U (1997) Markov games with incomplete information. *Math Meth Oper Res* 46:263–279
- Laraki R, Solan E (2013) Equilibrium in two-player nonzero-sum Dynkin games in continuous time. *Stochastics* 85:997–1014
- Laraki R, Solan E (2005) *SIAM Journal on Control and Optimization* 43:1913–1922
- Laraki R, Solan E, Vieille N (2005) Continuous time games of timing. *J Econ Theory* 120:206–238
- Levhari D, Mirman L (1980) The great fish war: an example using a dynamic Cournot-Nash solution. *Bell J Econ* 11(1):322–334
- Levy Y (2012) Stochastic games with information lag. *Games Econ Behav* 74:243–256
- Levy, Y. (2013a) Discounted stochastic games with no stationary Nash equilibrium: two examples, *Econometrica*, 81, 1973–2007.
- Levy Y (2013b) Continuous-time stochastic games of fixed duration. *Dyn Games Appl* 3:279–312
- Levy Y, McLennan A (2015) Corrigendum to “discounted stochastic games with no stationary Nash equilibrium: two examples”. *Econometrica* 83:1237–1252
- Lovo S, Tomala T (2015) Markov perfect equilibria in stochastic revision games, HEC Paris research paper no. ECO/SCD-2015-1093

- Maitra A, Sudderth W (1998) Finitely additive stochastic games with Borel measurable payoffs. *Int J Game Theory* 27:257–267
- Mamer JW (1987) Monotone stopping games. *J Appl Probab* 24:386–401
- Martin DA (1998) The determinacy of Blackwell games. *J Symb Log* 63:1565–1581
- Mashiach-Yaakovi A (2015) Correlated equilibria in stochastic games with Borel measurable payoffs. *Dyn Games Appl* 5:120–135
- Mertens JF (1987) Repeated games. In: *Proceedings of the international congress of mathematicians, Berkeley*, pp 1528–1577
- Mertens JF, Neyman A (1981) Stochastic games. *Int J Game Theory* 10:53–66
- Mertens JF, Parthasarathy T (1987) Equilibria for discounted stochastic games, CORE discussion paper no. 8750. Also published. In: Neyman A, Sorin S (eds) *Stochastic games and applications*, NATO science series. Kluwer, Dordrecht, pp 131–172
- Milman E (2006) Approachable sets of vector payoffs in stochastic games. *Games Econ Behav* 56:135–147
- Morimoto H (1986) Non zero-sum discrete parameter stochastic games with stopping times. *Probab Theory Relat Fields* 72:155–160
- Neyman, Sorin (2010) *Int J Game Theory* 39:29–52
- Neyman A (2008) Existence of optimal strategies in Markov games with incomplete information. *Int J Game Theory* 37:581–596
- Neyman A (2012) Continuous-time stochastic games. Discussion paper, Center for the Study of Rationality, Jerusalem, #616
- Neyman A (2013) Stochastic games with short-stage duration. *Dyn Games Appl* 3:236–278
- Neyman A, Sorin S (2010) Repeated games with public uncertain duration process. *Int J Game Theory* 39:29–52
- Nowak AS (1985) Existence of equilibrium stationary strategies in discounted Noncooperative stochastic games with uncountable state space. *J Optim Theory Appl* 45:591–620
- Nowak AS (2003a) Zero-sum stochastic games with Borel state spaces. In: Neyman A, Sorin S (eds) *Stochastic games and applications*, NATO science series. Kluwer, Dordrecht, pp 77–91
- Nowak AS (2003b) *N*-person stochastic games: extensions of the finite state space case and correlation. In: Neyman A, Sorin S (eds) *Stochastic games and applications*, NATO science series. Kluwer, Dordrecht, pp 93–106
- Nowak AS (2003c) On a new class of nonzero-sum discounted stochastic games having stationary Nash equilibrium points. *Int J Game Theory* 32:121–132
- Nowak AS, Raghavan TES (1992) Existence of stationary correlated equilibria with symmetric information for discounted stochastic games. *Math Oper Res* 17:519–526
- Ohtsubo Y (1987) A non zero-sum extension of Dynkin's stopping problem. *Math Oper Res* 12:277–296
- Ohtsubo Y (1991) On a discrete-time non-zero-sum Dynkin problem with monotonicity. *J Appl Probab* 28:466–472
- Papadimitriou C (2007) The complexity of computing equilibria. In: Nisan N, Roughgarden T, Tardos E, Vazirani VV (eds) *Algorithmic game theory*. Cambridge University Press, Cambridge
- Peski M, Wiseman T (2015) A folk theorem for stochastic games with infrequent state changes. *Theor Econ* 10:131–173
- Phelan C, Stacchetti E (2001) Sequential equilibria in a Ramsey tax model. *Econometrica* 69:1491–1518
- Purves RA, Sudderth WD (2011) Perfect information games with upper Semicontinuous payoffs. *Math Oper Res* 36:468473
- Raghavan TES, Syed Z (2002) Computing stationary Nash equilibria of undiscounted single-Controller stochastic games. *Math Oper Res* 27:384–400
- Raghavan TES, Syed Z (2003) A policy improvement type algorithm for solving zero-sum two-Person stochastic games of perfect information. *Math Program Ser A* 95:513–532
- Renault J (2006) The value of Markov chain games with lack of information on one side. *Math Oper Res* 31:490–512
- Renault J (2011) Uniform value in dynamic programming. *J Eur Math Soc* 13:309–330
- Renault J (2012) The value of repeated games with an informed controller. *Math Oper Res* 37:154–179
- Rosenberg D, Solan E, Vieille N (2001) Stopping games with randomized strategies. *Probab Theory Relat Fields* 119:433–451
- Rosenberg D, Solan E, Vieille N (2002) Blackwell optimality in Markov decision processes with partial observation. *Ann Stat* 30:1178–1193
- Rosenberg D, Solan E, Vieille N (2003) The maxmin of stochastic games with imperfect monitoring. *Int J Game Theory* 32:133–150
- Rosenberg D, Solan E, Vieille N (2004) Stochastic games with a single controller and incomplete information. *SIAM J Control Optim* 43:86–110
- Rosenberg D, Solan E, Vieille N (2009) Protocol with no acknowledgement. *Oper Res* 57:905–915
- Rubinstein A (1979) Equilibrium in supergames with the overtaking criterion. *J Econ Theory* 47:153–177
- Sagduyu YE, Ephremides A (2003) Power control and rate adaptation as stochastic games for random access. In: *Proceedings of the 42nd IEEE conference on decision and control*, vol 4, pp 4202–4207
- Savani R, von Stengel B (2004) Exponentially many steps for finding a Nash equilibrium in a bimatrix game. In: *Proceedings of the 45th annual IEEE symposium on foundations of computer science*, 2004, pp 258–267
- Shapley LS (1953) Stochastic games. *Proc Natl Acad Sci U S A* 39:1095–1100
- Simon RS (2003) The structure of non-zero-sum stochastic games. *Adv Appl Math* 38:1–26
- Simon RS (2007) The structure of non-zero-sum stochastic games. *Adv Appl Math* 38:1–26
- Simon RS (2012) A topological approach to Quitting games. *Math Oper Res* 37:180–195
- Solan E (1998) Discounted stochastic games. *Math Oper Res* 23:1010–1021
- Solan E (1999) Three-Person absorbing games. *Math Oper Res* 24:669–698

- Solan E (2001) Characterization of correlated equilibria in stochastic games. *Int J Game Theory* 30:259–277
- Solan E, Shmaya S (2004) Two-player non-zero-sum stopping games in discrete time. *Ann Probab* 32:2733–2764
- Solan E, Vieille N (2001) Quitting games. *Math Oper Res* 26:265–285
- Solan E, Vieille N (2002) Correlated equilibrium in stochastic games. *Games Econ Behav* 38:362–399
- Solan E, Vieille N (2010) Computing uniform optimal strategies in two-player stochastic games. *Econ Theor* 42:237–253. Special issue on Equilibrium Computation.
- Solan E, Vohra R (2002) Correlated equilibrium payoffs and public signalling in absorbing games. *Int J Game Theory* 31:91–122
- Sorin S (1984) Big match with lack of information on one side (part 1). *Int J Game Theory* 13:201–255
- Sorin S (1985) Big match with lack of information on one side (part 2). *Int J Game Theory* 14:173–204
- Sorin S (1986) Asymptotic properties of a non-zero-sum stochastic games. *Int J Game Theory* 15:101–107
- Sorin S, Vignal G (2015) Reversibility and oscillations in zero-sum discounted stochastic games. *J Dyn Games Appl* 2:103–115
- Sorin S, Zamir S (1991) Big match with lack of information on one side (Part 3). In: Raghavan TES et al (eds) *Stochastic games and related topics*. Kluwer, Dordrecht, pp 101–112
- Takahashi M (1964) Stochastic games with infinitely many strategies. *J Sci Hiroshima Univ Ser A-I* 26:123–134
- Thuijsman F, Raghavan TES (1997) Perfect information stochastic games and related classes. *Int J Game Theory* 26:403–408
- Vieille N (2000a) Equilibrium in 2-person stochastic games I: a reduction. *Israel J Math* 119:55–91
- Vieille N (2000b) Equilibrium in 2-person stochastic games II: the case of recursive games. *Israel J Math* 119:93–126
- Vignal G (2013) A zero-sum stochastic game with compact action sets and no asymptotic value. *Dyn Games Appl* 3:172–186
- Vrieze OJ, Thuijsman F (1989) On equilibria in repeated games with absorbing states. *Int J Game Theory* 18:293–310
- Vrieze OJ, Tijs SH (1982) Fictitious play applied to sequences of games and discounted stochastic games. *Int J Game Theory* 12:71–85
- Whitt W (1980) Representation and approximation of non-cooperative sequential games. *SIAM J Control Optim* 18:33–48
- Yasuda M (1985) On a randomized strategy in Neveu's stopping problem. *Stoch Process Appl* 21:159–166
- Zachrisson LE (1964) Markov games. In: Dresher M, Shapley LS, Tucker AW (eds) *Advances in game theory*. Princeton University Press, Princeton, pp 211–253
- Ziliotto B (2016) Zero-sum repeated games: counterexamples to the existence of the asymptotic value and the conjecture $\maxmin = \liminf_n$. *Ann Probab* 44:1107–1133
- ## Books and Reviews
- Başar T, Olsder GJ (1995) *Dynamic noncooperative game theory*. Academic, London
- Başar T, Zaccour G (2017) *Handbook of dynamic game theory*. Springer, Birkhäuser Basel
- Forges F (2007) Correlated equilibrium and communication in games. *This Encyclopedia*
- Friedman A (1971) *Differential games, pure and applied mathematics*, vol 25. Wiley, New York
- Jaśkiewicz A, Nowak AS (2017a) Zero-sum stochastic games. In: Başar T, Zaccour G (eds) *Handbook of dynamic game theory*. Springer International Publishing AG. Springer International Publishing. <https://doi.org/10.1007/978-3-319-27335-8-8-1>
- Jaśkiewicz A, Nowak AS (2017b) Non-zero-sum stochastic games. In: Başar T, Zaccour G (eds) *Handbook of dynamic game theory*. Springer International Publishing AG. Springer International Publishing. <https://doi.org/10.1007/978-3-319-27335-8-33-1>
- Kifer Y (2013) Dynkin's games and Israeli options. *ISRN Probab Stat*. Article ID:856458
- Laraki R, Sorin S (2015) Advances in zero-sum dynamic games. In: Aumann RJ, Hart S (eds) *Handbook of game theory with economic applications*, vol 4. Elsevier, North-Holland, pp 27–94
- Maitra AP, Sudderth WD (1996) *Discrete gambling and stochastic games*. Springer, Berlin
- Mertens JF (2002) Stochastic games. In: Aumann RJ, Hart S (eds) *Handbook of game theory with economic applications*, vol 3. Elsevier, North-Holland, pp 1809–1832
- Mertens JF, Sorin S, Zamir S (2015) *Repeated games*. Cambridge University Press, New York
- Neveu J (1975) *Discrete-parameter martingales*. North-Holland, Amsterdam
- Neyman A, Sorin S (2003) *Stochastic games and applications*, NATO science Series. Kluwer, Dordrecht
- Nisan N, Roughgarden T, Tardos E, Vazirani VV (eds) (2007) *Algorithmic game theory*. Cambridge University Press, Cambridge/New York
- Puterman ML (1994) *Markov decision processes: discrete stochastic dynamic programming*. Wiley, New York
- Raghavan TES, Shapley LS (1991) *Stochastic games and related topics: in honor of Professor L.S. Shapley*. Springer Science+Business Media Dordrecht
- Ramachandran KM, Tsokos CP (2012) *Stochastic differential games: theory and application*, Atlantis studies in probability and statistics, vol 2. Atlantis Press, Paris
- Sorin S (2002) A first course on zero-sum repeated games, *Mathématiques et applications*, vol 37. Springer, Berlin
- Sorin S (2003) Stochastic games with incomplete information. In: Neyman A, Sorin S (eds) *Stochastic games and applications*, NATO science series. Kluwer, Dordrecht, pp 375–395
- Vieille N (2002) Stochastic games: recent results. In: Aumann RJ, Hart S (eds) *Handbook of game theory with economic applications*, vol 3. Elsevier, Amsterdam, pp 1833–1850

Signaling Games

Joel Sobel

Department of Economics, University of California, San Diego, CA, USA

Article Outline

Glossary

Definition of the Subject

Introduction

The Model

Equilibrium

The Basic Model

Cheap Talk

Verifiable Information

Communication About Intentions

Applications

Future Directions

Bibliography

Glossary

Babbling equilibrium An equilibrium in which the sender's strategy is independent of type and the receiver's strategy is independent of signal.

Behavior strategy A strategy for an extensive-form game that specifies the probability of taking each action at each information set.

Behavioral type A player in a game who is constrained to follow a given strategy.

Cheap-talk game A signaling game in which players' preferences do not depend directly on signals.

Condition D1 An equilibrium refinement that requires out-of-equilibrium beliefs to be supported on types that have the most to gain from deviating from a fixed equilibrium.

Divinity An equilibrium refinement that requires out-of-equilibrium beliefs to place

relatively more weight on types that gain more from deviating from a fixed equilibrium.

Equilibrium outcome The probability distribution over terminal nodes determined by an equilibrium strategy in a game.

Handicap principle The idea that animals communicate fitness through observable characteristics that reduce fitness.

Incomplete information game A game in which players lack information about the strategy sets or payoff functions of their opponents.

Intuitive criterion An equilibrium refinement that requires out-of-equilibrium beliefs to place zero weight on types that can never gain from deviating from a fixed equilibrium outcome.

Nash equilibrium A strategy profile in a game in which each player's strategy is a best response to the equilibrium strategies of the other players.

Neologism-proof equilibrium An equilibrium that admits no self-signaling set.

Pooling equilibrium A signaling-game equilibrium in which each all sender types send the same signal with probability one.

Receiver In a signaling game, the uninformed player.

Self-signaling set A set of types C with the property that precisely types in the set C gain from inducing the best response to C relative to a fixed equilibrium.

Sender In a signaling game, the informed agent.

Separating equilibrium A signaling-game equilibrium in which sender types send signals from disjoint subsets of the set of available signals.

Signaling game A two-player game of incomplete information in which one player is informed and the other is not. The informed player's strategy is a type-contingent message and the uninformed player's strategy is a message-contingent action.

Single-crossing condition A condition that guarantees that indifference curves from a given family of preferences cross at most once.

Spence-Mirrlees condition A differential condition that orders the slopes of level sets of a function.

Standard signaling game A signaling game in which strategy sets and payoff functions satisfy monotonicity properties.

Type In an incomplete information game, a variable that summarizes private information.

Verifiable information game A signaling game with the property that each type has a signal that can only be sent by that type.

Definition of the Subject

Signaling games refer narrowly to a class of two-player games of incomplete information in which one player is informed and the other is not. The informed player's strategy set consists of signals contingent on information, and the uninformed player's strategy set consists of actions contingent on signals. More generally, a signaling game includes any strategic setting in which players can use the actions of their opponents to make inferences about hidden information. The earliest work on signaling games was Spence's (1974) model of educational signaling and Zahari's (1975) model of signaling by animals. During the 1980s researchers developed the formal model and identified conditions that permitted the selection of unique equilibrium outcomes in leading models.

Introduction

The framed degree in your doctor's office, the celebrity endorsement of a popular cosmetic, and the telephone message from an old friend are all signals. The signals are potentially valuable because they allow you to infer useful information. These signals are indirect and require interpretation. They may be subject to manipulation. The doctor's diploma tells you something about the doctor's qualifications, but knowing where and when the doctor studied does not prove that she is a good doctor. The endorsement identifies the product with a particular lifestyle, but what works for the celebrity may not work for you. Besides, the celebrity was probably paid to endorse the product and may not even use it. The phone message may tell you how to get in touch with your friend, but is unlikely to contain all of the

information you need to find him – or to evaluate whether you will meet to discuss old times or to be asked a favor. While the examples all involve signaling, the nature of the signaling is different. The doctor faces large penalties for misrepresenting her credentials. She is not required to display all of her diplomas, but it is reasonable to assume that degrees are not forged. The celebrity endorsement is costly – certainly to the manufacturer who pays for the celebrity's services and possibly to the celebrity himself, whose reputation may suffer if the product works badly. It is reasonable to assume that it is easier to obtain an endorsement of a good product, but there are also good reasons to be skeptical about the claims. In contrast, although a dishonest or misleading message may lead to a bad outcome, leaving a message is not expensive and the content of the message is not constrained by your friend's information. The theory of signaling games is a useful way to describe the essential features of all three examples.

Opportunities to send and evaluate signals arise in many common natural and economic settings. In the canonical example (due to Spence 1974), a high-ability worker invests in education to distinguish herself from less skilled workers. The potential employer observes educational attainment, but not innate skill, and infers that a better educated worker is more highly skilled and pays a higher wage. To make this story work, there must be a reason that low-ability workers do not get the education expected of a more highly skilled worker and hence obtain a higher wage. This property follows from an assumption that the higher the ability of the worker, the easier it is for her to produce a higher signal.

The same argument appears in many applications. For example, a low-risk driver will purchase a lower cost, partial insurance contract, leaving the riskier driver to pay a higher rate for full insurance (Rothschild and Stiglitz 1976 or Wilson 1977). A firm that is able to produce high-quality goods signals this by offering a warranty for the goods sold (Grossman 1981) or advertising extensively. A strong deer grows extra large antlers to show that it can survive with this handicap and to signal its fitness to potential mates (Zahavi 1975).

Game theory provides a formal language to study how one should send and interpret signals

in strategic environments. This entry reviews the basic theory of signaling and discusses some applications. It does not discuss related models of screening. Kreps and Sobel (1994) and Riley (2001) review both signaling and screening.

Section “[The Model](#)” describes the basic model. Section “[Equilibrium](#)” defines equilibrium for the basic model. Section “[The Basic Model](#)” limits attention to a special class of signaling game. I give conditions sufficient for the existence of equilibria in which the informed agent’s signal fully reveals her private information and argue that one equilibrium of this kind is prominent. The next three sections study different signaling games. Section “[Cheap Talk](#)” discusses models of costless communication. Section “[Verifiable Information](#)” discusses the implications of the assumptions that some information is verifiable. Section “[Communication About Intentions](#)” briefly discusses the possibility of signaling intentions rather than private information. Section “[Applications](#)” describes some applications and extensions of the basic model. Section “[Future Directions](#)” speculates on directions for future research.

The Model

This section describes the basic signaling game. There are two players, called S (for sender) and R (for receiver). S knows the value of some random variable t whose support is a given set T . t is called the **type** of S . The prior beliefs of R are given by a probability distribution $\pi(\cdot)$ over T ; these beliefs are common knowledge. When T is finite, $\pi(t)$ is the prior probability that the sender’s type is t . When T is uncountably infinite, $\pi(\cdot)$ is a density function. Player S learns t and sends to R a signal s drawn from some set M . Player R receives this signal and then takes an action a drawn from a set A . (It is possible to allow A to depend on s and M to depend on t .) This ends the game: The payoff to i is given by a function $u^i: T \times M \times A \rightarrow \mathbb{R}$.

This canonical game captures the essential features of the classic applications of market signaling. In the labor market signaling story due to Spence (1974), a worker wishes to signal his ability to a potential employer. The worker has information about ability that the employer lacks. Direct

communication about ability is not possible, but the worker can acquire education. The employer can observe the worker’s level of education and use this to form a judgment about the worker’s true level of ability. In this application, S is a worker; R represents a potential employer (or a competitive labor market); t is the worker’s productivity; s is her level of education; and a is her wage.

Equilibrium

Defining Nash equilibrium for the basic signaling game is completely straightforward when T , M , and A are finite sets. In this case a behavior strategy for S is a function $\mu: T \times M \rightarrow [0, 1]$ such that $\sum_{s \in M} \mu(t, s) = 1$ for all t . $\mu(t, s)$ is the probability that sender-type t sends the signal s . A behavior strategy for R is a function $\alpha: M \times A \rightarrow [0, 1]$ where $\sum_{a \in A} \alpha(s, a) = 1$ for all s . $\alpha(s, a)$ is the probability that R takes action a following the signal s .

Proposition 1 Behavior strategies (α^*, μ^*) form a Nash equilibrium if and only if for all $t \in T$

$$\begin{aligned} \mu(t, s) > 0 \text{ implies } \sum_{a \in A} U^S(t, s, a) \alpha(s, a) \\ = \max_{s' \in S} \sum_{a \in A} U^S(t, s', a) \alpha(s', a) \end{aligned} \quad (1)$$

and for each $s \in M$ such that $\sum_{t \in T} \mu(t, s) \pi(t) > 0$,

$$\begin{aligned} \alpha(s, a) > 0 \text{ implies } \sum_{t \in T} U^R(t, s, a) \beta(t, a) \\ = \max_{a' \in A} \sum_{t \in T} U^R(t, s, a') \beta(t, a'), \end{aligned} \quad (2)$$

where

$$\beta(t, s) = \frac{\mu(t, s) \pi(t)}{\sum_{t' \in T} \mu(t', s) \pi(t')}. \quad (3)$$

Condition (1) states that the S places positive probability only on signals that maximize expected utility. This condition guarantees that

S responds optimally to R 's strategy. Condition (2) states that R places positive probability only on actions that maximize expected utility, where the expectation is taken with respect to the distribution $\beta(\cdot, s)$ following the signal s . Condition (3) states that $\beta(\cdot, s)$ accurately reflects the pattern of play. It requires that R 's beliefs be determined using S 's strategy and the prior distribution whenever possible. Equilibrium refinements also require that R has beliefs following signals s that satisfy

$$\sum_{t \in T} \mu(t, s) \pi(t) = 0, \quad (4)$$

that is, those signals that are sent with probability zero in equilibrium. Specifically, sequential equilibrium permits $\beta(\cdot, s)$ to be an arbitrary distribution when Eq. 4 holds, but requires that Eq. 2 holds even for these values of s . This restriction rules out equilibria in which certain signals are not sent because the receiver responds to the signal with an action that is dominated.

The ability to signal creates the possibility that R will be able to draw inferences about S 's type from the signal. Whether he is able to do so is a property of the equilibrium. It is useful to define two extreme cases.

Definition 1 An equilibrium (α^*, μ^*) is called a separating equilibrium if each type t sends different signals. That is, M can be partitioned into sets M_t such that for each t , $\sum_{s \in M_t} \mu(t, s) = 1$. An equilibrium (α^*, μ^*) is called a pooling equilibrium if there is a single signal s^* that is sent by all types with probability one.

In a separating equilibrium, R can infer S 's private information completely. In a pooling equilibrium, R learns nothing from the sender's signal. This definition excludes other possible situations. For example, all sender types can randomize uniformly over a set of two or more signals. In this case, the receiver will be able to draw no inference beyond the prior from a signal received in equilibrium. More interesting is the possibility that the equilibrium will be partially revealing, with some, but not all of the sender types sending common signal. It is not

difficult to construct pooling equilibria for the basic signaling game. Take the labor market model and assume S sends the message s^* with probability one and that the receiver responds to s^* with his best response to the prior distribution and to all other messages with the best response to the belief that t is the least skilled agent. Provided that the least skilled agent prefers to send s^* to sending the cheapest alternative signal, this is a Nash equilibrium.

The Basic Model

The separating equilibrium is a benchmark outcome for signaling games. When a separating equilibrium exists, then it is possible for the sender to share her information fully with the receiver in spite of having a potential conflict of interest.

Existence of separating equilibria typically requires a systematic relationship between types and signals. An appropriate condition, commonly referred to as the single-crossing condition, plays a prominent role in signaling games and in models of asymmetric information more generally.

In this section I limit attention to a special class of signaling game in which there is a monotonic relationship between types and signals. In these models, separating equilibria typically exist.

I begin by stating the assumption in the environment most commonly seen in applications. Assume that the sets T , M , and A are all real intervals.

Definition 2 $U^S(\cdot)$ satisfies the single-crossing condition if $U^S(t, s, a) \leq U^S(t, s', a')$ for $s' > s$ implies that $U^S(t', s, a) < U^S(t', s', a')$ for all $t' > t$.

In a typical application, $U^S(\cdot)$ is strictly decreasing in its second argument (the signal) and increasing in its third argument (R 's response) for all types. Consequently indifference curves are well defined in $M \times A$ for all t . The single-crossing condition states that indifference curves of different sender types cross once. If a lower type is indifferent between two signal-action pairs, then a higher type strictly prefers to send the higher signal. In this way, the single-crossing condition guarantees that higher types send weakly higher signals in equilibrium.

Note two generalizations of Definition 2. First, the assumption that the domain of $U^S(\cdot)$ is the product of intervals can be replaced by the assumption that these sets are partially ordered. In this case, weak and strict order replaces the weak and strict inequalities comparing types and actions in the statement of the definition. Second, it is sometimes necessary to extend the definition to mixed strategies. In this case, the ordering of A induces a partial ordering of distributions of A through first-order stochastic dominance.

When one thinks of the single-crossing condition geometrically, it is apparent that it implies a ranking of the slopes of the indifference curves of the sender. Suppose that $U^S(\cdot)$ is smooth, strictly increasing in actions and strictly decreasing in signals so that indifference curves are well defined for each t . For fixed t , an indifference curve is a set of the form $\{(s, a): U^S(t, s, a) \equiv c\}$ for some constant c . Denote the function by $\bar{a}(s; t)$, so that $U^S(t, s, \bar{a}(s; t)) \equiv c$. It follows that the slope of the indifference curve of a type t sender is

$$\bar{a}_1(s; t) = -\frac{U_2^S(t, s, a)}{U_3^S(t, s, a)}, \quad (5)$$

where $\bar{a}_1(s; t)$ is the partial derivative of $\bar{a}(s; t)$ with respect to the first argument, and $U_k^S(\cdot)$ denotes the partial derivative of $U^S(\cdot)$ with respect to its k th argument. Under these conditions, the single-crossing condition is implied by the requirement that the right-hand side of Eq. 5 is decreasing in t . The differentiable version of the single-crossing condition is often referred to as the Spence-Mirrlees condition. Milgrom and Shannon (1994) contain general definitions of the single-crossing and Spence-Mirrlees conditions, and Edlin and Shannon (1998) provide a precise statement of when the conditions are equivalent.

To provide a simple construction of a separating equilibrium, limit attention to a standard signaling game in which the following conditions hold:

1. $T = \{0, \dots, K\}$ is finite.
2. A and M are real intervals.
3. Utility functions are continuous in action and signal.

4. $U^S(\cdot)$ is strictly increasing in action and strictly decreasing in signal.
5. The single-crossing property holds.
6. The receiver's best-response function is uniquely defined, independent of the signal, and strictly increasing in t so that it can be written $BR(t)$.
7. There exists $\bar{s} \in M$ such that $U^S(K, \bar{s}, BR(K)) < U^S(K, s_0^*, BR(0))$.

Conditions (1) and (2) simplify exposition, but otherwise are not necessary. It is important that T , M , and A be partially ordered so that some kind of single-crossing condition applies. Conditions (4)–(6) impose a monotone structure on the problem so that higher types are more able to send high signals and that higher types induce higher (and uniformly more attractive) actions. These conditions imply that in equilibrium, higher types will necessarily send weakly higher signals. Condition (7) is a boundary condition that makes sending high signals unattractive. It states that the highest type of sender would prefer to be treated like the lowest type rather than use the signal $\bar{s}$. These properties hold in many standard applications. Condition (6) would be satisfied if $U^R(t, s, a) = -(a - t)^2$.

Separating Equilibrium

To illustrate these ideas, consider a construction of a separating equilibrium.

Proposition 2 The standard signaling game has a separating equilibrium.

One can prove the proposition by constructing a possible equilibrium path and confirming that the path can be part of a separating equilibrium.

1. $t = 0$ selects the signal s_0^* that maximizes $U^S(0, s, BR(0))$.
2. Suppose that s_i^* have been specified for $i = 0, \dots, k - 1$ and let $U^*(i) = U^S(i, s_i^*, BR(i))$.

. Define s_k^* to solve

$$\max U^S(k, s, BR(k)) \text{ subject to } U^S(k - 1, s, BR(k)) \leq U^*(k - 1).$$

Provided that the optimization problems in Steps 1 and 2 have solutions, the process

inductively produces a signaling strategy for the sender and a response rule for the receiver defined on $\{s_0^*, \dots, s_K^*\}$. When $BR(\cdot)$ is strictly increasing, the single-crossing condition implies that the signaling strategy is strictly increasing. To complete the description of strategies, assume that the receiver takes the action $BR(k)$ in response to signals in the interval $[s_k, s_{k+1})$, $BR(0)$ for $s < s_0^*$, and $BR(K)$ for $s > s_K^*$. By the definition of the best-response function, the receiver is best responding to the sender's strategy. When the boundary condition fails, a separating equilibrium need not exist, but when M is compact, one can follow the construction above to obtain an equilibrium in which the lowest types separate and higher types pool at the maximum signal in M (see Cho and Sobel 1990 for details).

In the construction, the equilibrium involves inefficient levels of signaling. When $U^S(\cdot)$ is decreasing in the signal, all but the lowest type of sender must make a wasteful expenditure in the signal in order to avoid being treated as having a lower quality. The result that expenditures on signals are greater than the levels optimal in a full-information model continue to hold when $U^S(\cdot)$ is not monotonic in the signal. The sender inevitably does no better in a separating equilibrium than she would do if R had full information about t . Indeed, all but the lowest type will do strictly worse in standard signaling games. On the other hand, the equilibrium constructed above has a constrained efficiency property: Of all separating equilibria, it is Pareto dominant from the standpoint of S . To confirm this claim, argue inductively that in any separating equilibrium, if t_j sends the signal s_j , then $s_j \geq s_j^*$, with equality only if all types $i < j$ send s_i^* with probability one.

Mailath (1987) provides a similar construction when T is a real interval. In this case, the Spence-Mirrlees formulation of the single-crossing condition plays an important role, and the equilibrium is a solution to a differential equation.

Multiple Equilibria and Selection

Section “[Equilibrium](#)” ended with the construction of a pooling equilibrium. A careful reconsideration of the argument reveals that there typically

are many pooling equilibrium outcomes. One can construct a potential pooling outcome by assuming that all sender types send the same signal and that the receiver best responds to this common signal and responds to all other signals with the least attractive action. Under the standard monotonicity assumptions, this strategy profile will be an equilibrium if the lowest sender type prefers pooling to sending the cheapest available out-of-equilibrium message. Section “[Separating Equilibrium](#)” ended with the construction of a separating equilibrium. There are also typically many separating equilibrium outcomes. Assume that types $t = 0, \dots, r - 1$ send signals s_t^* , type r sends $\tilde{s}_r > s_r^*$, and subsequent signals $\tilde{s}_t$ for $t > r$ solve

$$\begin{aligned} \max_s U^S(t, s, BR(t)), \text{ subject to } U^S(t - 1, s, BR(t)) \\ \leq U(t - 1, \tilde{s}_{t-1}, BR(t - 1)). \end{aligned}$$

In both of these cases, the multiplicity is typically profound, with a continuum of distinct equilibrium outcomes (when M is an interval). The multiplicity of equilibria means that, without refinement, equilibrium theory provides few clear predictions beyond the observation that the lowest type of sender receives at least $U^*(0)$, the payoff it would receive under complete information, and the fact that the equilibrium signaling function is weakly increasing in the sender's type. The first property is a consequence of the monotonicity of S 's payoff in a and of R 's best response function. The second property is a consequence of the single-crossing condition.

This section describes techniques that refine the set of equilibria. Refinement arguments that guarantee existence and select unique outcomes for standard signaling games rely on the Kohlberg-Mertens (1986) notion of strategic stability. The complete theory of strategic stability is only available for finite games. Consequently the literature applies weaker versions of strategic stability that are defined more easily for large games. Banks and Sobel (1987), Cho and Kreps (1987), and Cho and Sobel (1990) introduce these arguments.

Multiple equilibria arise in signaling games because Nash equilibrium does not constrain the

receiver's response to signals sent with zero probability in equilibrium. Specifying that R 's response to these unsent signals is unattractive leads to the largest set of equilibrium outcomes. (In standard signaling games, S 's preferences over actions do not depend on type, so the least attractive action is well defined). The equilibrium set shrinks if one restricts the meaning of unsent signals. An effective restriction is condition D1, introduced in Cho and Kreps (1987). This condition is less restrictive than the notion of universal divinity introduced by Banks and Sobel (1987), which in finite games is less restrictive than Kohlberg and Mertens's notion of strategic stability.

Given an equilibrium (α^*, μ^*) , let $U^*(t)$ be the equilibrium expected payoff of a type t sender and let $D(t, s) = \{a : u(t, s, a) \geq U^*(t)\}$ be the set of pure-strategy responses to s that lead to payoffs at least as great as the equilibrium payoff for player t . Given a collection of sets, $\{X(t)\}_{t \in T}$, $X(t^*)$ is maximal if it is not a proper subset of any $X(t)$.

Definition 3 Behavior strategies (α^*, μ^*) together with beliefs β^* satisfy D1 if for any unsent message s , $\beta^*(\cdot, s)$ is supported on those t for which $D(t, s)$ is maximal.

In standard signaling games, $D(t, s)$ is an interval: all actions greater than or equal to a particular action will be attractive relative to the equilibrium. Hence, these sets are nested. If $D(t, s)$ is not maximal, then there is another type t' that is "more likely to deviate" in the sense that there exists out-of-equilibrium responses that are attractive to t' but not t . Condition D1 requires that the receiver place no weight on type t making a deviation in this case. Notice if $D(t, s)$ is empty for all t , then D1 does not restrict beliefs given s (and any choice of action will support the putative equilibrium). Condition D1 is strong. One can imagine weaker restrictions. The intuitive condition (Cho and Kreps 1987) requires that $\beta^*(t, s) = 0$ when $D(t, s) = \emptyset$ and at least one other $D(t', s)$ is nonempty. Divinity (Banks and Sobel 1987) requires that if $D(t, s)$ is strictly contained in $D(t', s)$, then $\beta^*(t', s)/\beta^*(t, s) \geq \pi(t')/\pi(t)$, so that the relative probability of the types more likely to deviate increases.

Proposition 3 The standard signaling game has a unique separating equilibrium outcome that satisfies condition D1.

In standard signaling games, the only equilibrium outcome that satisfies condition D1 is the separating outcome described in the previous section. Details of the argument appear in Cho and Sobel. The argument relies on two insights. First, types cannot be pooled in equilibrium because slightly higher signals will be interpreted as coming from the highest type in the pool. Second, in any separating equilibrium in which a sender type fails to solve Step 2, deviation to a slightly lower signal will not lower R 's beliefs.

The refinement argument is powerful and the separating outcome selected receives prominent attention in the literature. It is worth pointing out that the outcome has one unreasonable property. The separating outcome described above depends only on the support of types and not on the details of the distribution. Further, all types but the lowest type must make inefficient (compared to the full-information case) investments in signal in order to distinguish themselves from lower types. The efficient separating equilibrium for a sequence of games in which the probability of the lowest type converges to zero does not converge to the separating equilibrium of the game in which the probability of the lowest type is zero. In the special case of only two types, the (efficient) pooling outcome may be a more plausible outcome when the probability of the lower type shrinks to zero. Grossman and Perry (1987) and Mailath et al. (1993) introduce equilibrium refinements that select the pooling equilibrium in this setting. These concepts share many of the same motivations of the refinements introduced by Banks and Sobel and Cho and Kreps. They are qualitatively different from the intuitive criterion, divinity, and condition D1, because they are not based on dominance arguments and lack general existence properties.

Cheap Talk

Models in which preferences satisfy the single-crossing property are central in the literature, but

the assumption is not appropriate in some interesting settings. This section describes an extreme case in which there is no direct cost of signaling.

In general, a cheap-talk model is a signaling model in which $U^i(t, s, a)$ is independent of s for all (t, a) . Two facts about this model are immediate. First, if equilibrium exists, then there always exists an equilibrium in which no information is communicated. To construct this “babbling” equilibrium, assume that $\beta(t, s)$ is equal to the prior independent of the signal s . R ’s best response will be to take an action that is optimal conditional only on his prior information. Hence, R ’s action can be taken to be constant. In this case, it is also a best response for S to send a signal that is independent of type, which makes $\beta(t, s)$ the appropriate beliefs. Hence, even if the interests of S and R are identical, so that there are strong incentives to communicate, there is a possibility of complete communication breakdown.

Second, it is clear that nontrivial communication requires that different types of S have different preferences over R ’s actions. If it is the case that whenever some type t prefers action a to action a' , then so do all other types, then (ruling out indifference¹) it must be the case that in equilibrium the receiver takes only one action with positive probability. To see this, note that otherwise one type of sender is not selecting a best response. The second observation shows that cheap talk is not effective in games, like the standard labor-market story, in which the sender’s preferences are monotonic in the action of the receiver. With cheap communication, the potential employee in the labor market will always select a signal that leads to the highest possible wage and consequently, in equilibrium, all types of workers will receive the same wage.

A Simple Cheap-Talk Game

There are natural settings in which cheap talk is meaningful in equilibrium. To describe examples, I follow the development of Crawford and Sobel (1982). (Green and Stokey (2007) independently

introduced a similar game in an article circulated in 1981.) In this entry, A and T are the unit interval and M can be taken to be the unit interval without loss of generality. The sender’s private information or type, t , is drawn from a differentiable probability distribution function, $F(\cdot)$, with density $f(\cdot)$, supported on $[0, 1]$. S and R have twice continuously differentiable von Neumann-Morgenstern utility functions $U^i(t, a)$ that are strictly concave in a and have a strictly positive mixed partial derivative. For $i = R, S$, $a^i(t)$ denotes the unique solution to $\max_a U^i(t, a)$ and further assume that $a^S(t) > a^R(t)$ for all t . (The assumptions on $U^i(\cdot)$ guarantee that $U^i(\cdot)$ is well defined and strictly increasing.)

In this model, the interests of the sender and receiver are partially aligned because both would like to take a higher action with a higher t . The interests are different because S would always like the action to be a bit higher than R ’s ideal action. In a typical application, t represents the ideal action for R , such as the appropriate expenditure on a public project. Both R and S want actual expenditure to be close to the target value, but S has a bias in favor of additional expenditure.

For $0 \leq t' < t'' \leq 1$, let $\bar{a}(t', t'')$ be the unique solution to $\max_a \int_{t'}^{t''} U^R(a, t) dF(t)$. By convention, $\bar{a}(t, t) = a^R(t)$.

Without loss of generality, limit attention to pure-strategy equilibria. The concavity assumption guarantees that R ’s best responses will be unique, so R will not randomize in equilibrium. An equilibrium with strategies (μ^*, α^*) induces action a if $\{t : \alpha^*(\mu^*(t)) = a\}$ has positive prior probability. Crawford and Sobel (1982) characterize equilibrium outcomes.

Proposition 4 There exists a positive integer N^* such that for every integer N with $1 \leq N \leq N^*$, there exists at least one equilibrium in which the set of induced actions has cardinality N , and moreover, there is no equilibrium which induces more than N^* actions. An equilibrium can be characterized by a partition of the set of types, $t(N) = (t_0(N), \dots, t_N(N))$ with $0 = t_0(N) < t_1(N) < \dots < t_N(N) = 1$, and signals $m_i, i = 1, \dots, N$, such that for all $i = 1, \dots, N - 1$

¹Chakraborty and Rick Harbaugh (2010) construct informative equilibria in a game in which the sender’s preferences are independent of type.

$$U^S(t_i, \bar{a}(t_i, t_{i+1})) - U^S(t_i, \bar{a}(t_{i-1}, t_i)) = 0, \quad (6)$$

$$\mu(t) = m_i \quad \text{for } t \in (t_{i-1}, t_i], \quad (7)$$

and

$$\alpha(m_i) = \bar{a}(t_{i-1}, t_i). \quad (8)$$

Furthermore, essentially all equilibrium outcomes can be described in this way.

In an equilibrium, adjacent types pool together and send a common message. Condition (6) states that sender types on the boundary of a partition element are indifferent between pooling with types immediately below or immediately above. Condition (7) states that types in a common element of the partition send the same message. Condition (8) states that R best responds to the information in S 's message.

Crawford and Sobel make another monotonicity assumption, which they call condition (M). (M) is satisfied in leading examples and implies that there is a unique equilibrium partition for each $N = 1, \dots, N^*$, the ex ante equilibrium expected utility for both S and R is increasing in N , and N^* increases if the preferences of S and R become more aligned. These conclusions provide justification for the view that with fixed preferences “more” communication (in the sense of more actions induced) is better for both players and that the closer the interests of the players are, the greater the possibilities for communication.

As in the case of models with costly signaling, there are multiple equilibria in the cheap-talk model. The multiplicity is qualitatively different. Costly signaling models have a continuum of Nash equilibrium outcomes. Cheap-talk models have only finitely many. Refinements that impose restrictions on off-the-equilibrium path beliefs work well to identify a single outcome in costly signaling models. These refinements have no cutting power in cheap-talk models because any equilibrium distribution on type-action pairs can arise from signaling strategies in which all messages are sent with positive probability. To prove this claim, observe that if message m' is unused in equilibrium, while message m is used, then one can construct a new equilibrium in which

R interprets m' the same way as m and sender types previously sending m randomize equally between m and m' .

In the basic model messages take on meaning only through their use in an equilibrium. Unlike natural language, they have no external meaning. There have been several attempts to formalize the notion that messages have meanings that, if consistent with strategic aspects of the interaction, should be their interpretation inside the game. The first formulation of this idea is due to Farrell (1993).

Definition 4 Given an equilibrium with sender expected payoffs $U * (\cdot)$, the subset $G \subset T$ is self-signaling if $G = \{t : U^S(t, BR(G)) > U * (t)\}$.

That is, G is self-signaling if precisely the types in G gain by making a statement that induces the action that is a best response to the information that $t \in G$. (When $BR(t)$ is not single valued, it is necessary to refine the definition somewhat and permit the possibility that $U^S(t, BR(G)) = U * (t)$ for some t . See Mathews et al. (1991).) Farrell argues that the existence of a self-signaling set would destroy an equilibrium. If a subset G had available message that meant “my type is in G ,” then relative to the equilibrium R could infer that if he were to interpret the message literally, then it would be sent only by those types in G (and hence, the literal meaning would be accurate). With this motivation, Farrell proposes a refinement.

Definition 5 An equilibrium is neologism proof if there exist no self-signaling sets relative to the equilibrium.

Rabin (1990) argues convincingly that Farrell's definition rules out too many equilibrium outcomes. Indeed, for leading examples of the basic cheap-talk game, there are no neologism-proof equilibria. Specifically, in the Crawford-Sobel model in which S has a bias toward higher actions, there exist self-signaling sets of the form $[t, 1]$. On the other hand, Chen et al. (2007) demonstrate that the N^* -step equilibrium always satisfies the no incentive to separate (NITS) condition,

$$U^S(0, \alpha * (\mu * (0))) \geq U^S(0, \alpha^R(0)), \quad (9)$$

and that under condition (M) this is the only equilibrium that satisfies condition (9).

NITS states that the lowest type of sender prefers her equilibrium payoff to the payoff she would receive if the receiver knew her type (and responded optimally). Kartik (2009) introduced and named this condition. The NITS condition can be shown to rule out equilibria that admit self-signaling sets of the form $[0, t]$. Chen (2011) and Kartik (2009) show that the condition holds in the limits of perturbed versions of the basic cheap-talk game.

Condition 9 holds automatically in any perfect Bayesian equilibrium of the standard signaling model. This follows because when R 's actions are monotonic in type and S 's preferences are monotonic in action, the worst outcome for S is to be viewed as the lowest type. This observation would not be true in Nash equilibrium, where it is possible for R to respond to an out-of-equilibrium message with an action $a < BR(0)$.

Variations on Cheap Talk

In standard signaling models, there is typically an equilibrium that is fully revealing. This is not the case in the basic cheap-talk model. This leads to the question of whether it is possible to obtain more revelation in different environments.

One idea is to consider the possibility of signaling over many dimensions. Chakraborty and Harbaugh (2007) consider a model in which T and A are multidimensional. A special case of their model is one in which the components of T are independent draws from the same distribution and A involves taking a real-valued action for each component of T . If preferences are additively separable across types and actions, Chakraborty and Harbaugh provide conditions under which categorical information transmission, in which S transmits the order of the components of T , is credible in equilibrium even when it would not be possible to transmit information if the dimensions were treated in isolation. It may be credible for S to say " $t_1 > t_2$," even if she could not credibly provide information about the absolute value of either component of t .

Communication is nontrivial if some sender type strictly prefers to induce one equilibrium action over another. Nontrivial communication typically requires that different types have different preferences over outcomes. In standard signaling models, the heterogeneity arises because different sender types have different costs of sending messages. In cheap-talk models with one-dimensional actions, the heterogeneity arises if different sender types have different ideal actions. With multidimensional actions, heterogeneity could come simply from different sender types having different preferences over the relative importance of the different issues. Another simple variation is to assume the existence of more than one sender. In the two-sender game, nature picks t as before, both senders learn t and simultaneously send a message to the receiver, who makes a decision based on the two messages. The second sender has preferences that depend on type and the receiver's action, but not directly on the message sent. In this environment, assume that $M = T$, so that the set of available messages is equal to the type space (this is essentially without loss of generality). One can look for equilibria in which the senders report honestly. Denote by $a^*(t, t')$ R 's response to the pair of messages (t, t') . If an equilibrium in which both senders report honestly exists, then R 's response to identical messages is $a^*(t, t) = a^R(t)$. It must be the case that there exists a specification of $a^*(t, t')$ for $t \neq t'$ such that for all $i = 1$ and 2 and $t \neq t'$,

$$U^{S_i}(t, a^*(t, t)) \geq U^{S_i}(t, a^*(t, t')). \quad (10)$$

It is possible to satisfy condition (10) if the biases of the senders are small relative to the set of possible best responses. Krishna and Morgan (2001a) study a one-dimensional model of information transmission with two informed players. Ambrus and Takahashi (2008) and Battaglini (2002) provide conditions under which full revelation is possible when there are two informed players and possibly multiple dimensions of information.

In many circumstances, enriching the communication structure either by allowing more rounds of communication (Aumann and Hart 2003;

Forges 1990), mediation (Ben-Porath 2003), or exogenous uncertainty (Blume et al. 2007) enlarges the set of equilibrium outcomes.

Verifiable Information

Until now, the focus has been on situations in which the set of signals available does not depend on the true state. There are situations in which this assumption is not appropriate. There may be laws that ban false advertisement. The sender may be able to document details about the value of t . Models of this kind were first studied by Grossman (1981) and Milgrom (1981). For example, if t is the sender's skill at playing the piano, then if there is a piano available, the type t sender could demonstrate that she has skill at least as great as t (by performing at her true ability), but she may not be able to prove that her skill is no more than t (the receiver may think that she deliberately played the piano badly).

To model these possibilities, suppose that the set of possible messages is the set of all subsets of T . In this case, messages have "literal" meanings: When the sender uses the message $s = C \subset T$, this can be interpreted as a statement of the form: "my type is in C ." If the sender cannot lie, then $M(t)$ must be the set of subsets of T that contain t . If type t is verifiable, then $\{t\} \in M(t')$ if and only if $t' = t$. If there are no additional costs of sending signals, this model can be viewed as a variation of cheap-talk models in which the message space depends on t . In general, one can treat verifiable information models as a special case of the general signaling game in which the cost of sending certain signals is so large that these signals can be ruled out. Lying is impossible if $M(t) = \{C \in 2^T : t \in C\}$. In this setting, it is appropriate to require equilibria to be consistent with the signaling structure.

Definition 6 The equilibrium (σ^*, α^*) is rationalizable if

$$\begin{aligned} \alpha(C, a) > 0 & \text{ implies } \sum_{t \in T} U^R(t, s, a) \beta(t, s) \\ &= \max_{a' \in A} \sum_{t \in T} U^R(t, s, a') \beta(t, s), \end{aligned} \quad (11)$$

where $\beta(t, s) = 0$ if $t \notin C$.

Condition 11 requires that beliefs place positive probability only on types capable of sending the message "my type is an element of C ."

Proposition 5 Suppose that A and T are linearly ordered, that the receiver's best response function is increasing in type, and that all sender types prefer higher actions. If lying is not possible, then in any rationalizable equilibrium (α^*, μ^*) , $\alpha^*(s, BR(t)) = 1$ whenever $\mu^*(t, s > 0)$.

Grossman (1981) and Milgrom (1981) present versions of this proposition. Seidman and Winter (1997) generalize the result.

Provided that the receiver responds to the signal $\{t\}$ with $BR(t)$, each type can guarantee a payoff of $BR(t)$. On the other hand, if any type receives a payoff greater than $BR(t)$, then some higher type must be doing worse. Another way to make the same point is to notice that the highest type $\{\bar{t}\}$ has a weakly dominant strategy to reveal her type by announcing $\{\bar{t}\}$. Once this type is revealed, the next highest type will want to reveal herself and so on. Hence, verifiable information will be revealed voluntarily in an environment where cheap talk leads to no revealing and costly signaling will be compatible with full revelation, but at the cost of dissipative signaling.

The full-revelation result depends on the assumption that the sender and receiver share a linear ranking over the quality of information. Giovannini and Seidmann (2007) discuss more general settings in which the ability to provide verifiable information need not lead to full revelation.

Communication About Intentions

In a simple signaling game, signals permit uninformed agents to learn something about the private information of informed agents. Another possibility is to add a round of pre-play communication to a given game. Even if the game has complete information, there is the possibility that communication would serve to select equilibria or permit correlation that would otherwise be infeasible. Farrell and Rabin's (1996) review article discusses this literature in more detail.

Aumann (1990) argues that one cannot rely on pre-play communication to select a Pareto-efficient equilibrium. He considers a simple two-player game with Pareto-ranked equilibria and argues that no “cheap” pre-play signal would be credible.

Ben-Porath and Dekel (1992) show that adding a stage of “money burning” (a signal that reduces all future payoffs by the same amount) when combined with an equilibrium refinement can select equilibria in a complete information game. Although no money is burned in the selected equilibrium outcome, the potential to send costly signals creates dominance relationships that lead to a selection.

Vida (2006) synthesizes a literature that compares the set of equilibrium outcomes available when communication possibilities are added to a game to the theoretically larger set of equilibrium outcomes if there is a reliable mediator available to collect information and recommend actions to the players.

Applications

Economic Applications

There is an enormous literature that uses signaling models in applications. Riley’s (2001) survey contains extended discussion of some of the most important applications. What follows is a brief discussion of some central ideas.

In a simple signaling game, one informed agent sends a single signal to one uninformed decision-maker. This setting is general enough to illustrate many important aspects of signaling, but it is plainly limited. Interesting new issues arise if there are many informed agents, if there are many decision-makers, and if the interaction is repeated. Several of the models below add some or all of these novel features to the basic model.

Advertising

Advertisements are signals. Models similar to the standard model can explain situations in which higher levels of advertisement can lead consumers to believe the quality of the good is higher. In a separating equilibrium, advertising expenditures

fully reveal quality. As in all costly signaling models, it is not important that there be a direct relationship between quality and signal, it is only necessary that firms with higher quality have lower marginal costs of advertising. Hence, simply “burning money” or sending a signal that lowers utility by an amount independent of quality and response can be informative. The consumer may obtain full information in equilibrium, but someone must pay the cost of advertising. There are other situations where it is natural for the signal to be linked to the quality of the item. Models of verifiable information are appropriate in this case. When the assumptions of Proposition 5 hold, one would expect consumers to obtain all relevant information through disclosures without wasteful expenditures on signaling. Finally, cheap talk plays a role in some markets. One would expect costless communication to be informative in environments where heterogeneous consumers would like to identify the best product. Cheap talk can create more efficient matching of product to consumer. Here communication is free although in leading models separating equilibria do not exist.

Limit Pricing

Signaling models offer one explanation for the phenomenon of limit pricing. An incumbent firm has private information about its cost. Potential entrants use the pricing behavior of the firm to draw inferences about the incumbent’s cost, which determines profitability of entry. Milgrom and Roberts (1982a) construct an equilibrium in which the existence of incomplete information distorts prices: Relative to the full information model, the incumbent charges lower prices in order to signal that the market is relatively unprofitable. This behavior has the flavor of classical models of limit pricing, with one important qualification. In a separating equilibrium the entrant can infer the true cost of the incumbent, and therefore the low price charged by the incumbent firm fails to change the entry decision.

Bargaining

Several authors have proposed bargaining models with incomplete information to study the

existence and duration of strikes (Fudenberg and Tirole 1983; Sobel and Takahashi 1983). If a firm with private information about its profitability makes a take-it-or-leave-it offer to a union, then the strategic interaction is a simple signaling model in which the magnitude of the offer may serve as a signal of the firm's profitability. Firms with low profits are better able to make low wage offers to the union because the threat of a strike is less costly to a firm with low profits than one with high profits. Consequently settlement offers may reveal information. Natural extensions of this model permit counter offers. The variation of the model in which the uninformed agent makes offers and the uninformed agent accepts and rejects is formally almost identical to the canonical model of price discrimination by a durable-goods monopolist (Ausubel and Deneckere 1989; Gul et al. 1986).

Finance

Simple signaling arguments provide potential explanations for firms' choices of financial structure. Classic arguments due to Modigliani and Miller (1958) and imply that firms' profitability should not depend on their choice of capital structure. Hence, this theory cannot organize empirical regularities about firm's capital structure. The Modigliani-Miller theorem assumes that the firm's managers, shareholders, and potential shareholders all have access to the same information. An enormous literature assumes instead that the firm's managers have superior information and use corporate structure to signal profitability.

Leland and Pyle (1977) assume that insiders are risk averse so they would prefer to diversify their personal holdings rather than maintain large investments in their firm. The greater the value of diversification, the lower the quality of the firm. Hence, when insiders have superior information than investors, there will be an incentive for the insiders of highly profitable firms to maintain inefficiently large investments in their firm in order to signal profitability to investors.

Dividends are taxed twice under the US tax code, which raises the question of why firms would issue dividends when capital gains are taxed at a lower rate. A potential explanation for

this behavior comes from a model in which investors have imperfect information about the future profitability of the firm and profitable firms are more able than less profitable firms to distribute profits in the form of dividends (see Bhattacharya 1979).

Reputation

Dynamic models of incomplete information create the opportunity for the receiver to draw inferences about the sender's private information while engaging in an extended interaction. Kreps and Wilson (1982) and Milgrom and Roberts (1982b) provided the original treatments of reputation formation in games of incomplete information. Motivated by the limit pricing, their models examined the interaction of a single long-lived incumbent facing a sequence of potential entrants. The entrants lack information about the willingness of the incumbent to tolerate entry. Pricing decisions of the incumbent provide information to the entrants about the profitability of the market.

In these models, signals have implications for both current and future utility. The current cost is determined by the effect the signal has on current payoffs. In Kreps-Wilson and Milgrom-Roberts, this cost is the decrease in current profits associated with charging a low price. In other models (e.g., Morris 2001; Sobel 1985) the actual signal is costless, but it has immediate payoff implications because of the response it induces. Signals also have implications for future utility because inferences about the sender's private information will influence the behavior of the opponents in future periods. Adding concern for reputation to a signaling game will influence behavior, but whether it leads to more or less informative signaling depends on the application.

Signaling in Biology

Signaling is important in biology. In independent and almost contemporaneous work, Zahavi (1975) proposed a signaling model that shared the essential features of Spence's (1974) model of labor-market signaling. Zahavi observed that there are many examples in nature of animals' apparently excessive physical displays. It takes energy to produce colorful plumage, large antlers,

or loud cries. Having a large tail may actually make it harder for peacocks to flea predators. If a baby bird makes a loud sound to get his mother's attention, he may attract a dangerous predator. Zahavi argued that costly signals could play a role in sexual selection. In Zahavi's basic model, the sender is a male and the receiver is a female of the same species. Females who are able to mate with healthier males are more likely to have stronger children, but often the quality of a potential mate cannot be observed directly. Zahavi argued that if healthier males could produce visible displays more cheaply than less healthy males, then females would be induced to use the signals when deciding upon a mate. Displays may impose costs that "handicap" a signaler, but displays would persist when additional reproductive success compensates for their costs. Zahavi identifies a single-crossing condition as a necessary condition for the existence of costly signals.

The development of signaling in biology parallels that in economics, but there are important differences. Biology replaces the assumption of utility maximization and equilibrium with fitness maximization and evolutionary stability. That is, their models do not assume that animals consciously select their signal to maximize a payoff. Instead, the biological models assume that the process of natural selection will lead to strategy profiles in which mutant behavior has lower reproductive fitness than equilibrium behavior. This notion leads to static and dynamic solution concepts similar to Nash equilibrium and its refinements. Fitness in biological models depends on contributions from both parents. Consequently, a full treatment of signaling must take into account population genetics. Grafen (1990b) discusses these issues, and Grafen (1990a) and Siller (1998) provide further theoretical development of the handicap theory. Finally, one must be careful in interpreting heterogeneous quality in biological models. Natural selection should operate to eliminate the least fit genes in a population. To the extent that this arises, there is pressure for quality variation within a population to decrease over time. The existence of unobserved quality variations needed for signaling may be the result of relatively small variations about a population norm.

While most of the literature on signaling in biology focuses on the use of costly signals, there are also situations in which cheap talk is effective. A leading example is the "Sir Philip Sidney Game," originally developed by John Maynard Smith (1991) to illustrate the value of costly communication between a mother and child. The child has private information about its level of hunger, and the mother must decide to feed the child or keep the food for itself. Since the players are related, survival of one positively influences the fitness of the other. This creates a common interest needed for cheap-talk communication. There are two ways to model communication in this environment. The first is to assume that signaling is costly, with hungrier babies better able to communicate their hunger. This could be because the sound of a hungry baby is hard for sated babies to imitate, or it could be that crying for food increases the risk of predation and that this risk is relatively more dangerous to well-fed chicks than to starving ones (because the starving chicks have nothing to lose). This game has multiple equilibria in which signals fully reveal the state of the baby over a range of values (see Lachmann and Bergstrom 1998; Maynard Smith 1991). These papers look at a model in which both mother and child have private information. Alternatively, Bergstrom and Lachmann (1998) study a cheap-talk version of the game. Here there may be an equilibrium outcome in which the baby bird credibly signals whether or not he is hungry. Those who signal hunger get fed. The others do not. Well-fed baby birds may wish to signal that they are not hungry in order to permit the mother to keep food for herself. Such an equilibrium exists if the fraction of genes that mother and child share is large and the baby is already well fed.

Political Science

Signaling games have played an important role in formal models of political science. Banks (1991) reviews models of agenda control, political rhetoric, voting, and electoral competition. Several important models in this area are formally interesting because they violate the standard

assumptions frequently satisfied in economic models. I describe two such models in this subsection.

Banks (1990) studies a model of agenda setting in which the informed sender proposes a policy to a receiver (decision-maker), who can either accept or reject the proposal. If the proposal is accepted, it becomes the outcome. If not, then the outcome is a fallback policy. The fallback policy is known only to the sender. In this environment, the sender's strategy may convey information to the decision-maker. Signaling is costly, but, because the receiver's set of actions is binary, fully revealing equilibria need not exist. Refinements limit the set of predictions in this model to a class of outcomes in which only one proposal is accepted in equilibrium (and that this proposal is accepted with probability one), but there are typically a continuum of possible equilibrium outcomes.

Matthews (1989) develops a cheap-talk model of veto threats. There are two players, a Chooser (C), who plays the role of receiver, and a Proposer (P), who plays the role of sender. The players have preferences that are represented by single-peaked utility functions which depend on the real-valued outcome of the game and an ideal point. P 's ideal point is common knowledge. C 's ideal point is her private information, drawn from a prior distribution that has a smooth positive density on a compact interval. The game form is simple: C learns her type and then sends a cheap-talk signal to P , who responds with a proposal. C then either accepts or rejects the proposal. Accepted proposals become the outcome of the game. If C rejects the proposal, then the outcome is the status quo point.

As usual in cheap-talk games, this game has a babbling outcome in which C 's message contains no information and P makes a single, take-it-or-leave-it offer that is accepted with probability strictly between 0 and 1. Matthews shows there may be equilibria in which two outcomes are induced with positive probability (size-two equilibria), but size $n > 2$ (perfect Bayesian) equilibria never exist. In a size-two equilibrium, P offers his ideal outcome to those types of C whose message indicates that their ideal point is low; this offer is always accepted in equilibrium. If C indicated that

his ideal point is high, P makes a compromise offer that is sometimes accepted and sometimes rejected.

Future Directions

The most exciting developments in signaling games in the future are likely to come from interaction between economics and other disciplines.

Over the last 10 years the influence of behavioral economists has led the profession to rethink many of its fundamental models. An explosion of experimental studies has already influenced the interpretation of signaling models and has led to a reexamination of basic assumptions. There is evidence that economic actors lack the strategic sophistication assumed in equilibrium models. Further, economic agents may be motivated by more than their material well-being. Existing experimental evidence provides broad support for many of the qualitative predictions of the theory (Banks et al. 1994; Brandts and Holt 1992), but also suggests ways in which the theory may be inadequate.

The driving assumption of signaling models is that when informational asymmetries exist, senders will attempt to lie for strategic advantage and that sophisticated receivers will discount statements. These assumptions may be reconsidered in light of experimental evidence that some agents will behave honestly in spite of strategic incentives to lie. For example, Gneezy (2005) and Hurkens and Kartik (2009) present experimental evidence that some agents are reluctant to lie even when there is a financial gain from doing so. There is evidence from other disciplines that some agents are unwilling or unable to manipulate information for strategic advantage and that people may be well equipped to detect these manipulations in ways that are not captured in standard models (see, e.g., Ekman 2001 or Trivers 1971). Experimental evidence and, possibly, results from neuroscience may demonstrate that the standard assumption that some agents cannot manipulate information for their strategic advantage (or that other agents have ability to see through deception) will inform the development

of novel models of communication that include behavioral types. Several papers study the implications of including behavioral types into the standard paradigm. The reputation models of Kreps and Wilson (1982) and Milgrom and Roberts (1982a) are two early examples. Papers on communication by Chen (2011), Crawford (2003), Kartik (2005), and Olszewski (2004) are more recent examples. New developments in behavioral economics will inform future theoretical studies.

There is substantial interest in signaling in philosophy. Indeed, the philosopher David Lewis (2002) (first published in 1969) introduced signaling games prior to the contributions of Spence and Zahavi. Recently linguists have been paying more attention to game-theoretic ideas. Benz et al. (2005) collect recent work that attempts to formalize ideas from linguistic philosophy due to Grice (1991). While there have been a small number of contributions by economists in this area (Rubinstein 2000 and Sally 2005 are examples), there is likely to be more active interaction in the future.

Finally, future work may connect strategic aspects of communication to the actual structure of language. Blume (2000), Cucker et al. (2004), and Nowak and Krakauer (1999) present dramatically different models on how structured communication may result from learning processes. Synthesizing these approaches may lead to fundamental insights on how the ability to send and receive signals develops.

Acknowledgments I thank the Guggenheim Foundation, NSF, and the Secretaría de Estado de Universidades e Investigación del Ministerio de Educación y Ciencia (Spain) for financial support and Richard Brady, Kanako Goulding Hotta, and Jose Penalva for their comments. I am grateful to the Departament d'Economia i d'Història Econòmica and Institut d'Anàlisi Econòmica of the Universitat Autònoma de Barcelona for hospitality and administrative support.

Bibliography

Primary Literature

Ambrus A, Takahashi S (2008) Multi-sender cheap talk with restricted state space. *Theor Econ* 3(1):1–27

- Aumann RJ (1990) Nash equilibria and not self-enforcing. In: Gabszewicz JJ, Richard J-F, Wolsey LA (eds) *Economic decision making: games, econometrics and optimisation*. Elsevier, Amsterdam, pp 201–206
- Aumann R, Hart S (2003) Long cheap talk. *Econometrica* 71(6):1619–1660
- Ausubel LM, Deneckere RJ (1989) Reputation in bargaining and durable goods monopoly. *Econometrica* 57(3):511–531
- Banks JS (1990) Monopoly agenda control and asymmetric information. *Q J Econ* 105(2):445–464
- Banks JS (1991) *Signaling games in political science*. Routledge, Langhorne
- Banks JS, Sobel J (1987) Equilibrium selection in signaling games. *Econometrica* 55(3):647–661
- Banks J, Camerer C, Porter D (1994) An experimental analysis of nash refinements in signaling games. *Games Econ Behav* 6(1):1–31
- Battaglini M (2002) Multiple referrals and multidimensional cheap talk. *Econometrica* 70(4):1379–1401
- Ben-Porath E (2003) Cheap talk in games with incomplete information. *J Econ Theory* 108:45–71
- Ben-Porath E, Dekel E (1992) Signaling future actions and the potential for sacrifice. *J Econ Theory* 57:36–51
- Benz A, Jäger G, Van Rooij R (eds) (2005) *Game theory and pragmatics*. Palgrave MacMillan, Houndmills
- Bergstrom CT, Lachmann M (1998) Signalling among relatives. III. Talk is cheap. *Proc Natl Acad Sci U S A* 95:5100–5105
- Bhattacharya S (1979) Imperfect information, dividend policy, and the bird in the hand fallacy. *Bell J Econ* 10(1):259–270
- Blume A (2000) Coordination and learning with a partial language. *J Econ Theory* 95:1–36
- Blume A, Board O, Kawamura K (2007) Noisy talk. *Theor Econ* 2(4):396–440
- Brandts J, Holt CA (1992) An experimental test of equilibrium dominance in signaling games. *Am Econ Rev* 82(5):1350–1365
- Chakraborty A, Harbaugh R (2007) Comparative cheap talk. *J Econ Theory* 132(1):70–94
- Chakraborty A, Harbaugh R (2010) Persuasion by Cheap Talk. *Am Econ Rev* 100(5):2361–2382.
- Chen Y (2011) Perturbed communication games with honest senders and naive receivers. *J Econ Theory* 146(2):401–424.
- Chen Y, Kartik N, Sobel J (2007) On the robustness of informative cheap talk. Technical report, UCSD. *Econometrica* 76(1):117–136.
- Cho I-K, Kreps DM (1987) Signaling games and stable equilibria. *Q J Econ* 102(2):179–221
- Cho I-K, Sobel J (1990) Strategic stability and uniqueness in signaling games. *J Econ Theory* 50(2):381–413
- Crawford VP (2003) Lying for strategic advantage: rational and boundedly rational misrepresentation of intentions. *Am Econ Rev* 93(1):133–149
- Crawford VP, Sobel J (1982) Strategic information transmission. *Econometrica* 50(6):1431–1451
- Cucker F, Smale S, Zhou D-X (2004) Modeling language evolution. *Found Comput Math* 4(3):315–343

- Edlin AS, Shannon C (1998) Strict single-crossing and the strict spence-mirrlees condition: a comment on monotone comparative statics. *Econometrica* 66(6): 1417–1425
- Ekman P (2001) Telling lies. W.W. Norton, New York
- Farrell J (1993) Meaning and credibility in cheap-talk games. *Games Econ Behav* 5(4):514–531
- Forges F (1990) Equilibria with communication in a job market example. *Q J Econ* 105(2):375–398
- Fudenberg D, Tirole J (1983) Sequential bargaining with incomplete information. *Rev Econ Stud* 50(2): 221–247
- Giovannoni F, Seidmann DJ (2007) Secrecy, two-sided bias and the value of evidence. *Games Econ Behav* 59(2):296–315
- Gneezy U (2005) Deception: the role of consequences. *Am Econ Rev* 95(1):384–394
- Grafen A (1990a) Biological signals as handicaps. *J Theor Biol* 144:517–546
- Grafen A (1990b) Sexual selection unhandicapped by the fisher process. *J Theor Biol* 144:473–516
- Green JR, Stokey NL (2007) A two-person game of information transmission. *J Econ Theory* 135(1):90–104
- Grice HP (1991) Studies in the way of words. Harvard University Press, Cambridge
- Grossman S (1981) The role of warranties and private disclosure about product quality. *J Law Econ* 24:461–483
- Grossman S, Perry M (1987) Perfect sequential equilibria. *J Econ Theory* 39:97–119
- Gul F, Sonnenschein H, Wilson R (1986) Foundations of dynamic monopoly and the coase conjecture. *J Econ Theory* 39:155–190
- Hurkens S, Kartik N (2009) Would I Lie to You: On Social Preferences and Lying Aversion. *Experimental Econ* 12(2):180–192.
- Kartik N (2009) Information transmission with almost-cheap talk. *Rev Econ Stud* 76(4):1359–1395.
- Kohlberg E, Mertens J-F (1986) On the strategic stability of equilibria. *Econometrica* 54(5):1003–1037
- Kreps DM, Sobel J (1994) Signalling. In: Aumann RJ, Hart S (eds) Handbook of game theory: with economics applications, vol 2, Handbooks in economics, No 11. Elsevier, Amsterdam, pp 849–868, chap 25
- Kreps DM, Wilson R (1982) Reputation and imperfect information. *J Econ Theory* 27:253–277
- Krishna V, Morgan J (2001a) A model of expertise. *Q J Econ* 116(2):747–775
- Lachmann M, Bergstrom CT (1998) Signalling among relatives. II. Beyond the tower of babel. *Theor Popul Biol* 54:146–160
- Leland HE, Pyle DH (1977) Informational asymmetries, financial structure, and financial intermediation. *J Finance* 32(2):371–387
- Lewis D (2002) Convention: a philosophical study. Blackwell, Oxford
- Mailath GJ (1987) Incentive compatibility in signaling games with a continuum of types. *Econometrica* 55(6):1349–1365
- Mailath GJ, Okuno-Fujiwara M, Postlewaite A (1993) On belief based refinements in signaling games. *J Econ Theory* 60(2):241–276
- Mathews SA, Okuno-Fujiwara M, Postlewaite A (1991) Refining cheap-talk equilibria. *J Econ Theory* 55(2):247–273
- Mathews SA (1989) Veto threats: rhetoric in a bargaining game. *Q J Econ* 104(2):347–369
- Maynard Smith J (1991) Honest signalling: the Philip Sidney game. *Anim Behav* 42:1034–1035
- Milgrom PR (1981) Good news and bad news: representation theorems. *Bell J Econ* 21:380–391
- Milgrom P, Roberts J (1982a) Limit pricing and entry under incomplete information: an equilibrium analysis. *Econometrica* 50(2):443–459
- Milgrom P, Roberts J (1982b) Predation, reputation, and entry deterrence. *J Econ Theory* 27:280–312
- Milgrom PR, Shannon C (1994) Monotone comparative statics. *Econometrica* 62(1):157–180
- Modigliani F, Miller MH (1958) The cost of capital, corporation finance and the theory of investment. *Am Econ Rev* 48(3):261–297
- Morris S (2001) Political correctness. *J Polit Econ* 109:231–265
- Nowak MA, Krakauer DC (1999) The evolution of language. *Proc Natl Acad Sci* 96(14):8028–8033
- Olszewski W (2004) Informal communication. *J Econ Theory* 117:180–200
- Rabin M (1990) Communication between rational agents. *J Econ Theory* 51:144–170
- Rabin M, Farrell J (1996) Cheap talk. *J Econ Perspect* 10(3):103–118
- Riley JG (2001) Silver signals: twenty-five years of screening and signaling. *J Econ Lit* 39(2):432–478
- Rothschild M, Stiglitz J (1976) Equilibrium in competitive insurance markets: an essay on the economics of imperfect information. *Q J Econ* 90(4):629–649
- Rubinstein A (2000) Economics and language. Cambridge University Press, New York
- Sally D (2005) Can I say, bobobo and mean, there's no such thing as cheap talk? *J Econ Behav Organ* 57(3):245–266
- Seidmann DJ, Winter E (1997) Strategic information transmission with verifiable messages. *Econometrica* 65(1):163–169
- Siller S (1998) A note on errors in grafen's strategic handicap models. *J Theor Biol* 195:413–417
- Sobel J (1985) A theory of credibility. *Rev Econ Stud* 52(4):557–573
- Sobel J, Takahashi I (1983) A multistage model of bargaining. *Rev Econ Stud* 50(3):411–426
- Spence AM (1974) Market signaling. Harvard University Press, Cambridge
- Trivers RL (1971) The evolution of reciprocal altruism. *Q Rev Biol* 46:35–58
- Vida P (2006) Long pre-play communication in games. Ph D thesis, Autonomous University of Barcelona
- Wilson C (1977) A model of insurance markets with incomplete information. *J Econ Theory* 16:167–207
- Zahavi A (1975) Mate selection- a selection for a handicap. *J Theor Biol* 53:205–214

Books and Reviews

- Admati A, Perry M (1987) Strategic delay in bargaining. *Rev Econ Stud* 54:345–364
- Austen-Smith D (1990) Information transmission in debate. *Am J Polit Sci* 34(1):124–152
- Battigalli P, Siniscalchi M (2002) Strong belief and forward induction reasoning. *J Econ Theory* 106(2):356–391
- Bernheim BD (1994) A theory of conformity. *J Polit Econ* 102(5):841–877
- Blume A, Kim Y-G, Sobel J (1993) Evolutionary stability in games of communication. *Games Econ Behav* 5:547–575
- Fudenberg D, Tirole J (1991) *Game theory*. MIT Press, Cambridge, MA
- Gibbons R (1992) *Game theory for applied economists*. Princeton University Press, Princeton
- Kartik N, Ottaviani M, Squintani F (2007) Credulity, lies, and costly talk. *J Econ Theory* 134(1):93–116
- Krishna V, Morgan J (2001b) A model of expertise. *Q J Econ* 116(2):747–775
- Lo P-Y (2006) Common knowledge of language and iterative admissibility in a sender-receiver game. Technical report. Brown University
- Manelli AM (1996) Cheap talk and sequential equilibria in signaling games. *Econometrica* 64(4):917–942
- Milgrom P, Roberts J (1986) Price and advertising signals of product quality. *J Polit Econ* 94(4):796–821
- Noldeke G, Samuelson L (1997) A dynamic model of equilibrium selection in signaling markets. *J Econ Theory* 73(1):118–156
- Noldeke G, Van Damme E (1990) Signalling in a dynamic labour market. *Rev Econ Stud* 57(1):1–23
- Ottaviani M, Sørensen PN (2006) Professional advice. *J Econ Theory* 126(1):120–142
- Rabin M (1994) A model of pre-game communication. *J Econ Theory* 63(2):370–391
- Ramey G (1996) D1 signaling equilibria with multiple signals and a continuum of types. *J Econ Theory* 69(2):508–531
- Rasmusen EB (2006) *Games and information: an introduction to game theory*, 4th edn. Blackwell, New York
- Riley JG (1979) Informational equilibrium. *Econometrica* 47:331–359
- Sobel J, Stole L, Zapater I (1990) Fixed-equilibrium rationalizability in signaling games. *J Econ Theory* 52(2):304–331
- Spence AM (1973) Job market signaling. *Q J Econ* 90:225–243
- Swinkels JM (1999) Education signalling with preemptive offers. *Rev Econ Stud* 66(4):949–970

Inspection Games

Rudolf Avenhaus¹ and Morton J. Canty²

¹Armed Forces University Munich, Neubiberg, Germany

²Institute for Chemistry and Dynamics of the Geosphere, Forschungszentrum Jülich, Jülich, Germany

Article Outline

Glossary

Definition

Introduction

Selected Inspection Models

Future Directions

Bibliography

Glossary

Deterrence In an inspection game, deterrence is said to be achieved by a Nash equilibrium in which the inspectee behaves legally, or in accordance with the agreed rule.

Extensive form The extensive form of a noncooperative game is a graphical representation which describes a succession of moves by different players, including chance moves, and which can handle quite intricate information patterns.

Inspector leadership Leadership in inspection games is a strategic concept by which, through persuasive announcement of her strategy, the inspector can achieve deterrence.

Mixed strategy A mixed strategy for a player in a noncooperative game is a probability distribution over that player's pure strategies.

Nash equilibrium A Nash equilibrium in a noncooperative game is a specification of strategies for all players with the property that no player has an incentive to deviate unilaterally from her specified strategy. A *solution* of a

noncooperative game is a Nash equilibrium which is either unique, or which, for some reason, has been selected among alternative equilibria.

Noncooperative game An n -person noncooperative game in *normal* or *strategic form* is a list of actions, called *pure strategies*, for each of n players, together with a rule for specifying each player's payoff (utility) when every player has chosen a specific action. Each player seeks to maximize her own payoff.

Saddle point A saddle point is a Nash equilibrium of a two-person zero-sum game. The *value* of the game is the (unique) equilibrium payoff to the first player.

Utility Utilities are sequences of numbers assigned to the outcomes of any strategy combination which mirror the order of preferences of each player and which fulfill the axioms of von Neumann and Morgenstern.

Verification Verification is the independent confirmation by an inspector of the information reported by an inspectee. It is used most commonly in the context of arms control and disarmament agreements.

Zero-sum game A zero-sum game is a noncooperative game in which the payoffs of all players sum to zero for any specific combination of pure strategies.

Definition

Inspection games deal with the problem faced by an *inspector* who is required to control the compliance of an *inspectee* to some legal or otherwise formal undertaking. One of the best examples of an inspector, in the institutional sense, is the International Atomic Energy Agency (IAEA) which, under a United Nations mandate, verifies the activities of all States – inspectees – signatory to the Nuclear Weapons Non-proliferation Treaty. An inspection regime of this kind is a true conflict situation, even if the inspectee voluntarily submits

to inspection (in multinational or bilateral treaties this is invariably the case), because the *raison d'être* of any control authority must be the assumption that the inspected party has a real incentive to violate its commitments. The primary objective of the inspector is to *deter* the inspectee from illegal behavior or, barring this, to catch him out. It is thus natural that quantitative models of inspections should be non-cooperative games with at least two players, inspector and inspectee (s). This survey will be limited to just one inspectee, that is, we shall restrict ourselves to two-person non-cooperative inspection games.

Inspection games should be distinguished from related topics such as quality control or the prevention of random accidents, for which there are no adversaries that act strategically, or from search-and-destroy problems. The salient feature of an inspection game is that an inspector tries to prevent an inspectee from behaving illegally in terms of some commitment. The inspectee might, for example, decide not to violate, so that there is nothing to search for. In fact, deterrence is generally the inspector's highest priority. Nevertheless, a sharp distinction between, e.g., inspection games and quality control models cannot be made in all cases, as we will see in Subsection "Illegal Production."

Introduction

Immediately after von Neumann and Morgenstern's pioneering book *Theory of Games and Economic Behavior* (von Neumann and Morgenstern 1947), Arms Control and Disarmament (ACD) inspections may have been analyzed game-theoretically as classified military research; this is not known for sure but may be inferred from papers published later. Non-classified work started vigorously in the early 1960s with analyses for the United States Arms Control and Disarmament Agency (ACDA). These dealt with very general ACD problems, and also with concrete problems of test ban treaty verification. In that context probably the first genuine inspection game in the open literature was the recursive game developed by Dresher (1962). Since it was

seminal for later work it is presented in some detail in Subsection "Customs and Smugglers."

A second phase of inspection game development started around 1968 in connection with the verification of the Treaty on the Non-Proliferation of Nuclear Weapons (NPT). There was no model for this verification system, therefore new principles and tools had to be developed and analyzed, see, e.g., (Bierlein 1968, 1969; Höpfinger 1974). Because of its importance for the further development of the whole discipline of inspection games, one of the major components of NPT verification measures, namely material accountancy, will be discussed in Subsection "Diversion of Nuclear Material."

In Economics game-theoretic work on Accounting and Auditing was begun in the late 1960s. A first survey was given by Borch (1990). Since that time papers and books have been published regularly but on a limited scale along similar lines, placing emphasis on auditing practice, see, e.g., (Cavasoglu and Raghunathan 2004; Cook et al. 1997; Wilks and Zimbelman 2004) provided an updated review of theoretical and empirical research. In economic models known as *principal-agent problems*, in which inspections of economic transactions raise the question of their most efficient design, game-theoretic methods have been applied, early surveys having been given by Baiman (1982), Kanodia (1985) and Dye (1986).

In the last decade, new models have been developed and analyzed again in the context of recent ACD verification developments, in particular the more stringent requirements on re-negotiated NPT verification measures (IAEA 1997). Whereas under the previous regime purely technical aspects like size of the nuclear fuel cycle and accuracy of the measurement systems were considered, now qualitative features such as behavior and intentions of States had to be taken into account. This required the introduction of State-specific utilities, for first analyses of this kind see, e.g., (Avenhaus and Canty 1996; Kilgour 1992). Independently of concrete applications there has been an ongoing interest of mathematicians in the refinement and generalization of existing models, a few examples will be given below.

Inspections cause conflicts in many real world situations. In economics, there are services of many kinds the fulfillment or payment of which has to be verified. For example one is concerned with the central problem of principal-agent relationships, where the principal, e.g., an employer, delegates work or responsibility to the agent, the employee, and chooses a payment schedule that best exploits the agent's self-interests. The agent, of course, behaves so as to maximize her own utility given the fee schedule proposed by the principal.

Environmental agreements obviously give rise to inspection problems, but these have not yet received as much attention from modelers as one might have expected (and as they might deserve). To date most methodological analyses of inspection games have been performed in the context of arms control and disarmament.

There exist previous reviews of inspection games with objectives somewhat different to those of this survey. Avenhaus et al. (1996) restrict discussion to arms control and disarmament, and emphasize the historical development. Avenhaus et al. (2002) stress the methodological, and in particular the mathematical aspects. Here we undertake a new approach to organize the material, one which gives more credit to the diversity of the applications and the techniques necessary for their solution. We focus on selected game theoretic inspection models which, together with their variants and generalizations, we hope will span the full range of the subject.

Selected Inspection Models

In the following, five inspection problems are chosen to illustrate applications of inspection games together with their analysis. They are complemented with discussion of some of their variants and generalizations. In the last illustration, the special role of the leadership concept in inspection models is emphasized.

Passenger Ticket Control

A commuter on the Munich subway is, consciously or unconsciously, involved in a non-

cooperative game each time he boards a train. He can buy a valid ticket or travel "schwarz," risking a fine if he is controlled (checked). In its edition of July 18th, 1996, the daily *Süddeutsche Zeitung* reported the complaint of the Munich City Treasurer to the effect that the deployment of ticket inspectors by the local transit authority (MMV) was not worthwhile, the collected fines paying for only about half the cost of the inspectors themselves.

From the game theorist's viewpoint there is obviously an optimum control intensity that would alleviate this problem: employing just a single inspector would encourage many violations and result in her collecting more than enough fines to pay for herself, but would clearly not be in the MMV's interests. Using an army of inspectors would ensure compliance, but, there now being no fines at all, would not finance the inspectors. The optimum must lie between these two extremes.

Solution

The problem can be formulated as a two-person game in *normal form* involving the MVV as player 1 and the transit passenger as player 2 (Avenhaus 1997). The pure strategies for the MVV are to control or not to control, whereas the passenger will decide whether or not to buy a valid ticket. Let f be the fare, $b > f$ the fine and $e < b$ the control costs per passenger, all in euros. The game's normal form (also called *bimatrix form*) is shown in Fig. 1.

In the figure, the pure strategies of player 1 (control/no control) are depicted as rows and those of player 2 (legal/illegal) as columns. The payoffs to player 1 for each pure strategy combination are shown in the lower left hand corners of the corresponding squares, those for player 2 in the upper right hand corners. (This simple formulation ignores MVV overhead costs and any material gain the passenger may have from his trip.)

A solution of the game will be a *Nash equilibrium* (1951), that is, a pair of strategies, called *equilibrium strategies*, with the property that neither player has an incentive to deviate unilaterally from his or her equilibrium strategy. Equivalently, the strategies are said to be *mutual best replies*. In

		←	
1 \ 2	legal q	illegal $1 - q$	↑
	control p	no control $1 - p$	
control p	$-f$ $f - e$	$-b$ $b - e$	↑
no control $1 - p$	$-f$ f	0 0	
		→	

Inspection Games, Fig. 1 Normal form of the two-person game between MVV (*player 1*) and passenger (*player 2*). The arrows indicate the preference directions for the two players, the horizontal arrows for player 2, the vertical arrows for player 1

the Figure the preference directions, i.e., the deviation incentives, are seen to be cyclical. This means that there can be no Nash equilibrium involving pure strategies. However the equilibrium concept can be generalized to involve *mixed strategies* which are probability distributions over the sets of pure strategies. In the present case, the MVV controls with some probability p and the passenger behaves legally with probability q . The expected payoffs to the two players are then given by

$$\begin{aligned} E_1(p, q) &= (f - e)pq + (b - e)p(1 - q) + f(1 - p) \\ \times qE_2(p, q) &= -fpq - bp(1 - q) \\ &\quad - f(1 - p)q. \end{aligned} \quad (1)$$

If we designate the mixed equilibrium strategies by p^* and q^* and the equilibrium payoffs as $E_i^* = E_i(p^*, q^*)$, $i = 1, 2$ then the conditions for Nash equilibrium are

$$\begin{aligned} E_1^* &\geq E_1(p, q^*) \quad \text{for all } p \in [0, 1] \\ E_2^* &\geq E_2(p, q^*) \quad \text{for all } q \in [0, 1]. \end{aligned} \quad (2)$$

For this situation the equilibrium strategies can be determined simply by requiring that each

protagonist be made indifferent with regard to his or her own mixed strategy, see, e.g., Morris (1994). One obtains immediately

$$\begin{aligned} p^* &= \frac{f}{b}, \quad E_1^* = f \left(1 - \frac{e}{b}\right), \\ q^* &= 1 - \frac{e}{b}, \quad E_2^* = -f. \end{aligned} \quad (3)$$

At equilibrium the passenger behaves illegally with positive probability $1 - q^* = e/b$, that is, deterrence is not possible. We will return to this issue in Subsection “[Sharing Common Pool Resources](#).” Nevertheless on average he enjoys the same payoff as he would receive by paying his fare every time, namely $-f$.

Remarks

The average control expenditure for the MVV is ep , whereas the mean profit from collection of fines is $bp(1 - q)$. The difference is $(e - b(1 - q))p$. If the passenger plays his equilibrium strategy as given by Eq. (3), then

$$(e - b(1 - q^*))p = 0 \quad (4)$$

for any control probability p . The control costs are thus exactly compensated by the collected fines. It might be mentioned that the actual figures for inspection probability and income per passenger for Munich approximately satisfy (3). We may speculate that the City Treasurer’s complaint arises from the fact that regular violators develop strategies to recognize and avoid inspectors. Other complications not taken into account in the model are variations in frequency, hour of day and dwelling time of passengers using the system.

There are many inspection problems which can be described with models equivalent or similar to the one presented here. Inspection of metered parking spaces provides another example. Control of the sharing of common pool resources, as discussed in the last inspection model, below, belongs to the same category.

Illegal Production

In treaties prohibiting the production, acquisition and/or proliferation of weapons of mass

destruction, one is often concerned with the misuse of ostensibly legitimate production facilities. A commercial chemical plant, for example, may be used for production of forbidden precursors of chemical weapons, or a uranium enrichment facility may illegally enrich its product to weapons-grade U-235. The consequences of non-detection by an inspecting authority can be dire, and the *timeliness* of control procedures – the interval between the onset of plant misuse and its detection – may be especially important.

To illustrate, consider the following simple model. At the end of some reference time interval, for instance a calendar year or a production campaign, a major inspection takes place at a facility, one which would detect prior illegal production with certainty. Additionally, a single *interim inspection* is carried out, timed at the inspector's discretion, which will likewise detect prior violation with certainty. The interim inspection is intended to enhance the timeliness of detection should illegal activity be underway. The inspector would like to know precisely when it should take place.

Solution

This example entails solution of a *zero-sum game on the unit square*. The onset of illegal production and the time of the interim inspection are chosen on the interval strategically by the respective protagonists, plant operator and inspector. We take the payoff to the former as the time to detection of illegal production, and to the latter to be the negative of that quantity.

Representing the reference time by the interval $[0, 1]$, the operator's so-called *payoff kernel* is

$$A(y, x) = \begin{cases} y - x & \text{for } x \leq y \\ 1 - x & \text{for } x \geq y, \end{cases} \quad (5)$$

here $x \in [0, 1]$ denotes a pure strategy for the operator, the onset of illegal production and similarly $y \in [0, 1]$ is a pure inspection strategy. Let the inspector's and operator's *mixed strategy distribution functions* be $G(y)$ and $F(x)$, respectively. $G(y)$ is the probability of an inspection taking place at time y or earlier, $F(x)$ the probability of illegal activity beginning at time x or earlier. The

Nash equilibrium conditions, or, since we are dealing with a zero-sum game, the *saddle point criteria* determining the equilibrium mixed strategies G^* and F^* , are given by

$$E(G^*, x) \leq E(G^*, F^*) \leq E(y, F^*) \quad \text{for all } x, y \in [0, 1]. \quad (6)$$

Since the final inspection is certain, clearly the operator shouldn't wait too long to act. Rather, in constructing his optimal probability distribution $F^*(x)$, he might plausibly choose x randomly on an interval $[0, b]$ with $b < 1$. Consequently the inspector will not act later than b either. Let us assume that she chooses her inspection time y according to the probability density function $g^*(y)$, where

$$\int_0^b g^*(y) dy = 1. \quad (7)$$

The expected payoff to the operator for diversion at time $x \in [0, b]$ is then

$$\begin{aligned} E(G^*, x) &= \int_0^x (1 - x) g^*(y) dy \\ &\quad + \int_x^b (y - x) g^*(y) dy \\ &= \int_0^x g^*(y) dy + \int_x^b y g^*(y) dy \\ &\quad - x \int_0^b g^*(y) dy. \end{aligned} \quad (8)$$

But if the operator randomizes across the interval $[0, b]$ as assumed, this payoff must be constant for all $x \in [0, b]$ and equal to the value of the game, i.e., the equilibrium payoff to the operator. If this were not the case for some x in the interval, the operator would not have included it in his mixed strategy F^* in the first place. Thus the derivative of Eq. (8) with respect to x must vanish. This gives immediately

$$g^*(y) = \frac{1}{1 - y} \quad (9)$$

and from Eq. (7), $b = 1 - 1/e$. The expected payoff to the operator and value of the game is then easily seen to be $E(G^*, x) = 1/e$ for all $x \in [0, b]$.

Getting the operator's optimal strategy F^* is a bit more subtle because it requires a so-called *atom* at $X = 0$, that is to say, a finite probability of starting illegal production at precisely the beginning of the interval, as well as the probability density $f^*(x)$ on the remaining half-open interval $[0, b]$. In terms of the distribution function $F^*(x)$ that characterizes this mixed strategy, the atom is $F^*(0)$ and $f^*(x)$ is the derivative of $F^*(x)$ on $[0, b]$. The operator's payoff for some inspection time $y \in [0, b]$ is

$$\begin{aligned}
 E(y, F^*) &= yF^*(0) + \int_0^y (y-x)f^*(x)dx \\
 &\quad + \int_y^b (1-x)f^*(x)dx \\
 &= yF^*(0) + y \int_0^y f^*(x)dx \\
 &\quad + \int_y^b f^*(x)dx - \int_0^b xf^*(x)dx \\
 &= (y-1)F^*(y) + F^*(b) \\
 &\quad - \int_0^b xf^*(x)dx. \tag{10}
 \end{aligned}$$

Arguing as before, if the inspector randomizes over the interval, this expression must be constant for all $y \in [0, b]$. This is true if $(y-1)F^*(y)$ is independent of y . The requirement that $F^*(b) = 1$ then leads to

$$F^*(x) = \frac{1}{e} \cdot \frac{1}{1-x} \tag{11}$$

For $x \in [0, b]$ and $F^*(x) = 1$ for $x > b$. The atom is $F^*(0) = 1/e$; and the construction is complete. Verifying that F^* and G^* satisfy Eq. (6) is straightforward.

Remarks

Owen (1968) discussed the existence of Nash equilibria for continuous games on the unit square and methods for their solution. The above result was first obtained by Diamond (1982), who gave a generalization to $k > 2$ inspections. Prior to Diamond's work, Derman (1961) treated a somewhat similar minimax inspection problem.

Both models were motivated by reliability control problems: production units have to be inspected regularly and the earlier a failure is detected, the less costly it is for the production facility owner. Of course the production unit is not acting strategically, but in order to be on the safe side a minimax approach was chosen, which was then generalized by Diamond to give a saddle point solution. Thus we see that an approach which originally was not an inspection game according to our definition in Section "Definition" turned out to become one, with interesting applications such as the one discussed above.

Rothenstein and Zamir (2002) extended Diamond's model with a single inspection to include errors of the first and second kind. Krieger (2008) considered a time-discrete variant of the model. All variants require that both the operator and inspector commit themselves before the reference period begins. Thus if in the solution in Subsection "Solution" the operator simply waits for the interim inspection and then violates, he will achieve an expected time to detection of

$$\int_0^b (1-x)f^*(x)dx = b = 1 - \frac{1}{e} > \frac{1}{e},$$

and the inspector's advantage will have evaporated. But of course $f^*(x)$ is not the inspector's equilibrium in such a sequential game. Prior commitment may be justified in some cases, but not in others. If there is no requirement for commitment, the operator may prefer to start his illegal action immediately, i.e., at the beginning of the reference period, or delay his decision until the first intermediate inspection. This situation has to be modeled as an *extensive form game*. Its time-continuous version, which also considered errors of the first and second kind, was studied by Avenhaus and Canty (2005). Surprisingly it turned out that an equilibrium strategy of the inspector is a pure strategy, contrary to the equilibrium strategies of the Diamond-type models.

Diversion of Nuclear Material

As already mentioned in the introduction, a large number of game theoretic models of inspection situations have been developed in the framework of IAEA verification activities. The basis of the

IAEA inspection system is the verification of the continuing presence of fissile material in the peaceful nuclear fuel cycle of the State under consideration (IAEA 1972), therefore statistical measurement errors and, consequently, statistical decision theory must be taken into account.

Over a single accounting period, typically 1 year, we define the *material flow* as the measured net transfer of fissile material across the facility boundary, consisting of inputs (receipts of raw material) R and outputs (shipments of purified product and waste) S . If the physical inventory within the facility at the start of the period was I_0 , then the *book inventory* at the end of the accounting period is defined as

$$B = I_0 + R - S = I_0 + Y, \quad (12)$$

where $Y = R - S$ is the net material flow into the facility.

At the end of the period a new physical inventory I_1 is taken and compared with the book inventory,

$$Z = B - I_1 = I_0 + Y - I_1. \quad (13)$$

This expression, which is a random variable as a consequence of measurement error, defines the *material balance statistic*. If there are no unknown losses or diversions of material, its expectation value is

$$E(Z) = E(I_0) + E(Y) - E(I_1) = 0 \quad (14)$$

from conservation of mass. The quantity Z is commonly referred to as MUF, meaning *material unaccounted for*. Its reliable determination forms the basis for the inspector's conclusion with respect to non-diversion.

We shall focus upon a single nuclear facility and an inventory period of 1 year and pose the question: can the taking of *additional interim inventories* improve the detection sensitivity of the overall accountancy procedure?

Solution

The additional inventories essentially define a series of shorter material balance periods, say n in all. At the beginning of the first balance period, the amount I_0 of material subject to control

is measured in the facility. Then, during the i th period, $i = 1, \dots, n$, some net measured amount Y_i of material enters the area. At the end of that period the amount of material, now I_i is again measured. The quantity

$$Z_i = I_{i-1} + Y_i - I_i, \quad i = 1 \dots n,$$

is the material balance test statistic for the i th inventory period. Under the *null hypothesis* that no material was diverted, its expected value is, as before,

$$E_0(Z_i) = 0, \quad i = 1 \dots n.$$

The *alternative hypothesis* is that material is diverted from the balance area according to some specific pattern. Thus

$$E_1(Z_i) = \mu_i, \quad i = 1 \dots n, \quad \sum_{i=1}^n \mu_i = \mu > 0,$$

where the amount μ_i diverted in the i th period may be positive, negative or nil, while μ , the total amount of material missing, is hypothesized to be positive.

For the purpose of determining the best test procedure we now define a two-person zero-sum game, wherein the set of strategies of the inspector is the set of all possible test procedures $\{\delta_\alpha\}$, i.e., significance thresholds, for fixed false alarm probability α . The set of strategies of the operator is the set of diversion patterns $\boldsymbol{\mu} = (\mu_1 \dots \mu_n)^T$, $\sum_i \mu_i = \mu$. The payoff to the inspector is the probability of detection $1 - \beta(\delta_\alpha, \boldsymbol{\mu})$, where $\beta(\delta_\alpha, \boldsymbol{\mu})$ is the *second kind error probability* (= non-detection probability). A solution of the game is any strategy pair $(\delta_\alpha^*, \boldsymbol{\mu}^*)$ which satisfies the saddle point conditions

$$1 - \beta(\delta_\alpha^*, \boldsymbol{\mu}) \geq 1 - \beta(\delta_\alpha^*, \boldsymbol{\mu}^*) \geq 1 - \beta(\delta_\alpha, \boldsymbol{\mu}^*) \\ \text{for any } \delta_\alpha, \boldsymbol{\mu}. \quad (15)$$

With the aid of the Lemma of Neyman and Pearson, one of the most fundamental theorems in statistical decision theory, see, e.g., Rohatgi (1976), we can derive the following solution first obtained by Avenhaus and Jaech (1981). Suppose

that Σ is the covariance matrix of the multivariate normally distributed random vector $\mathbf{Z} = (Z_1, Z_2, \dots, Z_n)^T$ and define $\mathbf{e} = (1, 1, \dots, 1)^T$. Then the equilibrium strategies are in fact given by

$$\boldsymbol{\mu}^* = \frac{\mu}{\mathbf{e}^T \cdot \sum \cdot \mathbf{e}} \cdot \sum \cdot \mathbf{e}, \quad (16)$$

and by the test δ_α^* characterized by the critical region

$$\{z | \mathbf{e}^T \cdot \mathbf{z} > k_\alpha\} \quad (17)$$

where k_α is determined by α . The value of the game, that is, the guaranteed probability of detection, is given by

$$1 - \beta(\delta_\alpha^*, \boldsymbol{\mu}^*) = \Phi\left(\frac{\mu}{(\mathbf{e}^T \cdot \sum \cdot \mathbf{e})^{\frac{1}{2}}} - U(1 - \alpha)\right), \quad (18)$$

where Φ is the normal distribution and U is its inverse.

We can demonstrate that this is the case as follows: Under the null hypothesis of legal behavior

$$\begin{aligned} 1 - \alpha &= \text{Prob}_0(\mathbf{Z}^T \cdot \mathbf{e} < k_\alpha) \\ &= \Phi\left(\frac{k_\alpha}{(\text{var}(\mathbf{Z}^T \cdot \mathbf{e}))^{\frac{1}{2}}}\right) \end{aligned}$$

and therefore

$$\frac{k_\alpha}{(\text{var}(\mathbf{Z}^T \cdot \mathbf{e}))^{\frac{1}{2}}} = U(1 - \alpha).$$

Thus the left hand side of Eq. (15) is fulfilled as equality:

$$\begin{aligned} 1 - \beta(\delta_\alpha^*, \boldsymbol{\mu}) &= \text{Prob}_1(\mathbf{Z}^T \cdot \mathbf{e} > k_\alpha) \\ &= 1 - \Phi\left(\frac{k_\alpha - E_1(\mathbf{Z}^T \cdot \mathbf{e})}{(\text{var}(\mathbf{Z}^T \cdot \mathbf{e}))^{\frac{1}{2}}}\right) \\ &= \Phi\left(\frac{\mu}{(\mathbf{e}^T \cdot \sum \cdot \mathbf{e})^{\frac{1}{2}}} - U(1 - \alpha)\right). \end{aligned}$$

As for the right hand side, the critical region which maximizes the detection probability for μ^*

and for fixed α is, according to the Lemma of Neyman and Pearson, given by

$$\left\{z \mid \frac{f_1(z)}{f_0(z)} > k'_\alpha\right\} = \left\{z \mid \boldsymbol{\mu}^{*T} \cdot \sum^{-1} \cdot \mathbf{z} > k'_\alpha\right\},$$

where $f_i(z)$ are the joint density functions under hypothesis $i, i = 0, 1$. But from Eq. (16)

$$\boldsymbol{\mu}^{*T} \cdot \sum^{-1} \cdot \mathbf{z} \propto \left(\sum \cdot \mathbf{e}\right)^T \cdot \sum^{-1} \cdot \mathbf{z} = \mathbf{e}^T \cdot \mathbf{z}.$$

Thus δ_α^* is indeed a best reply to $\boldsymbol{\mu}^*$ and the right hand inequality (15) is fulfilled as well.

Remarks

According to Eq. (17), the inspector's optimal test statistic is

$$\mathbf{e}^T \cdot \mathbf{Z} = \sum_{i=1}^n Z_i = I_0 + \sum_{i=1}^n Y_i - I_n,$$

which is just the overall material balance for the entire time period involved. All of the intermediate inventories $I_i, i = 1, \dots, n - 1$, are *ignored*. This gives a definitive answer to the question as to whether additional inventory taking can improve the sensitivity of the material balancing system for detecting diversion. The answer is no.

Satisfying as it may be from a decision theoretic point of view, this result ignores the aspect of detection time. Waiting 1 year or one complete production campaign before evaluating the overall material balance may be too long to meet timeliness constraints. Therefore, under the name *near real-time material accountancy*, test procedures have been discussed which indeed subdivide the year into several inventory periods (at the cost of reduced overall detection sensitivity, as was just explained). To date, it has not been possible to define or to solve satisfactorily a decision theoretic model which takes the critical time aspect into account.

The IAEA safeguards system is organized in such a way that the inspector compares the material balance data reported by the plant operators (via their national authorities) with her own findings and thereafter, if no discrepancies are found,

closes the material balance with the help of the operator's reported data. Thus, along with material accountancy, *data verification* comprises the second foundation of the IAEA safeguards system. Due to the possibility that data may be intentionally falsified to make the material balance appear to be correct, data verification again poses game theoretic problems.

Two kinds of sampling procedures have to be considered in this context. In case of identifying items or checking of seals, so-called *attribute sampling* procedures are used in which only sampling errors have to be minimized. This leads on the one hand to stratified sampling solutions similar to those found in the context of accounting and auditing. One of these solutions became well-known, at least in expert circles, under the name IAEA formula, see e.g., (Avenhaus and Canty 1989). On the other hand, in case of quantitative destructive and non-destructive verification measurements, statistical measurement errors can no longer be avoided, leading to consideration of *variable sampling* procedures. A decision problem arises in this instance, since discrepancies between reported and independently verified data can be caused either by measurement errors or by real and intentionally generated differences (data falsification). Stewart (1971) was the first to propose the so-called *D*-statistic for use in IAEA data verification. For one class of data consisting of N items, n of which are verified, the *D*-statistic is the sum of the differences of reported data X_j and independently measured data Y_j , extrapolated to the whole class population, i.e.,

$$D_1 = \frac{N}{n} \sum_{j=1}^n (X_j - Y_j).$$

For K classes of data (for instance one class for each component of a closed material balance) the *D*-statistic is given by

$$D_K \sum_{i=1}^K \frac{N_i}{n_i} \sum_{j=1}^{n_i} (X_{ij} - Y_{ij}).$$

These quantities then form the basis for the test procedure of the inspector, which goes along similar lines as outlined before: Two hypotheses have to be formulated which permit the determination

of significance thresholds for fixed false alarm probabilities and, from them, the associated detection probabilities.

Later on (Avenhaus and Canty 1996) it was proven, again using the saddle point criterion and the Lemma of Neyman and Pearson, that the use of the *D*-statistic is optimal for a "reasonable" class of data falsification strategies, and it was shown how the sample sizes can be determined such that they maximize the overall probability of detecting a given total falsification for a total given inspection effort.

Customs and Smugglers

In the previous examples evidence of violation (illegal production, diversion) is assumed to persist: the illegal action can be detected after the fact. There are of course inspection problems where this kind of model is not appropriate. Probably the first genuine inspection game in the open literature was a recursive zero-sum game developed by Dresher (1962) which treated the situation in which the violator can only be caught red-handed, that is, if the illegal act is actually in progress when an inspection occurs. It's not difficult to imagine real situations where this is the case, a much-discussed example being the customs-smuggler problem (Thomas and Nisgav 1976). In its simplest form, a smuggler has n nocturnal opportunities to bring his goods safely across a channel. The customs office, equipped with a patrol boat, would very much like to apprehend him, but budget restrictions require that the boat can only patrol on $m < n$ nights. If a patrol coincides with a crossing attempt, the smuggler will be apprehended with certainty. Moreover the smuggler observes all patrols that take place. All that being the case, one can ask how customs should time its patrols.

Solution

The game theoretic model developed by Dresher (1962) at the RAND Corporation mentioned above fits this situation rather well. It illustrates nicely the special character of sequential games, and has an elegant recursive solution. We summarize it here, see von Stengel (1991) for a more thorough discussion as well as some interesting variations on the same theme.

In Dresher's model there are n time periods, during each of which the inspector can decide whether or not to control the inspectee, using up one of a total of m inspections available to her if she does so. The inspectee knows at each stage the number of past inspections. He can violate at most once, but can also choose to behave legally. Detection occurs only if violation and inspection coincide in the same period. The conflicting interests of the two players are again modeled as a zero-sum game, that is, the inspectee's loss on detection is the negative of the inspector's gain. Legal action gives a payoff of nil to both players. The game is shown in reduced *extensive form*, i.e., as a decision tree, in Fig. 2.

The inspectee's *information set*, shown as an oval in the figure, encompasses both of her decision points. This is meant to imply that she doesn't know at which node she is situated when choosing her strategy.

The entries at the leaf nodes of the tree are the payoffs to the inspector. The value of the game prior to the first period is denoted $v(n, m)$. If the single violation occurs, the inspector achieves +1 if an inspection takes place, otherwise -1. In the latter case, the game proceeds trivially with the inspectee behaving legally (he has already violated) and the inspector inspecting or not, as she chooses. If the inspectee behaves legally, the continuation of the game has, by definition, value $v(n-1, m-1)$ to the inspector if she decided to control in the first period, otherwise value $v(n-1, m)$. These values are the corresponding payoffs to the inspector after the first period, thus giving the recursive form of the game tree shown. The game terminates after detected violation or after the n periods, in the latter

case with a payoff of either 0 (legal behavior) or -1 (illegal behavior) to the inspector.

The function $v(n, m)$ is subject to two boundary conditions. If there are no periods left, and no violation has occurred,

$$v(0, m) = 0, \quad m > 0. \quad (19)$$

If the inspector has no inspections left, then the inspectee is aware of this and can safely violate (and will do so, since his payoff is higher):

$$v(n, 0) = -1, \quad n > 0. \quad (20)$$

We shall now seek an equilibrium of the game in the domain of mixed strategies. Let $p(n, m) \in [0, 1]$ be the probability with which the inspector chooses to inspect in the first period. The equilibrium choice for $p(n, m)$ makes the inspectee indifferent to legal or illegal behavior, so that he receives the same payoff $-v(n, m)$ given by

$$\begin{aligned} v(n, m) &= p(n, m)v(n-1, m-1) \\ &+ (1-p(n, m))v(n-1, m) \quad (\text{legal}) \\ v(n, m) &= p(n, m) \times (+1) + (1-p(n, m))(-1) \quad (\text{illegal}). \end{aligned}$$


In a similar way the inspector chooses her probability $q(n, m)$ for violation at the first stage so as to make the inspectee indifferent as to control or no control, leading to

$$\begin{aligned} q(n, m)v(n-1, m-1) + (1-q(n, m))1 \\ = q(n, m)v(n-1, m) + (1-q(n, m))(-1). \end{aligned}$$

Solving these three equations for $p(n, m)$, $q(n, m)$ and $v(n, m)$ we obtain

Inspection Games,

Fig. 2 Dresher's game in reduced extensive form



$$p(n, m) = \frac{v(n-1, m) + 1}{v(n-1, m) + 2 - v(n-1, m-1)}, m < n \quad (21)$$

and

$$q(n, m) = \frac{2}{v(n-1, m) + 2 - v(n-1, m-1)} \quad (22)$$

$$v(n, m) = \frac{v(n-1, m) + v(n-1, m-1)}{v(n-1, m) + 2 - v(n-1, m-1)}. \quad (23)$$

Equation 23 along with the two boundary conditions (19) and (20) determine $v(n, m)$ and therefore $p(n, m)$ and $q(n, m)$ uniquely. The explicit solution is given as follows: Define

$$t(n, m) = \sum_{i=1}^m \binom{n}{i}, \quad (24)$$

where $\binom{n}{i}$ denotes the binomial coefficient. For $0 < m < n$ the value of the game and the equilibrium strategies for both players are

$$v(n, m) = -\frac{\binom{n-1}{m}}{t(n, m)} \quad (25)$$

$$p(n, m) = \frac{t(n-1, m-1)}{t(n, m)} \quad (26)$$

$$q(n, m) = \frac{2}{\frac{\binom{n-2}{m}}{t(n-1, m)} - \frac{\binom{n-2}{m-2}}{t(n-1, m-1)} - 2} \quad (27)$$

It is not quite trivial to prove that this solution satisfies its determinants. In fact one has to use the following recursive relations for $t(n, m)$:

$$\begin{aligned} t(n-1, m) &= t(n-1, m-1) + \binom{n-1}{m} \\ t(n, m) &= t(n-1, m) + t(n-1, m-1). \end{aligned}$$

Remarks

Since the value of the game is negative, the inspectee will behave illegally with positive probability. (If at some stage in an actual play, the

inspector has as many inspections left as there are remaining opportunities, the inspectee will obviously behave legally.) Dresher's model uses abstract payoffs – utilities in the sense of von Neumann and Morgenstern (1947) – and therefore the zero-sum assumption becomes questionable. In some circumstances it can be argued that detected illegal action is, compared to legal action, *worse for both players*, inspector and inspectee. However, the non-zero-sum version of the game is not much more difficult to analyze and gives no additional structural insight.

Two technical remarks on Dresher's inspection game:

First of all, it works so well only because the concept of *behavioral strategies*, introduced by Kuhn (1953), can be applied. Kuhn showed that in a game in extensive form with *perfect recall* (i.e., a game in which both players remember their previous moves) mixed strategies can be replaced by sequences of behavioral strategies. These are probabilities which define the actions of both players at all information sets of the game. In fact, $p(n, m)$ and $q(n, m)$ defined above are behavioral strategies.

Second, the recursive approach requires that at each decision node there follow a *subgame* (i.e., a branch of the tree) which by definition may not be connected by a joint information set to other subgames. This, however is *not* the case in Dresher's game: if the control doesn't take place, the inspector doesn't know whether the inspectee has already behaved illegally and thus whether a true recursive subgame still exists. Fortunately this doesn't matter, since if the inspectee did in fact behave illegally the game is de facto over. Whatever strategy the inspector chooses subsequently will not affect her payoff.

As indicated initially, variants and extensions of Dresher's original model probably represent the largest class of inspection games formulated and analyzed so far, and new variants are still being published, see, e.g., (Hozaki et al. 2006; Pavlovic 2002). As also mentioned above, Dresher's zero sum assumption may be questioned – detected violation may still be worse for the inspector than no violation at all and, thus, negative for both players. Höpfinger (1971) was the first to introduce generalized


utilities and to determine the equilibria of the resulting recursive game. von Stengel (1991) considered the case that the inspectee intends to violate more than once, however only once at a given stage. He was able to solve the recursive game for at most k violations on n stages, $k < n$, under the assumption that the inspector is fully informed after each stage whether or not there is a violation, even when no inspection takes place, see also (Ferguson and Melolidakis 1998; Sakaguchi 1994). A third category of extensions is associated with the customs-smuggler problem, see, e.g., (Baston and Bostock 1991; Garnaev 1991; Goutal et al. 1997), where customs and the smuggler have more possibilities, e.g., customs has more than one boat, and the smugglers alternative routes. Finally, statistical errors of the first and second kind were considered, firstly by Maschler (1967) and later by Rinderle (1996). However, due to the complexity of the models and the many additional assumptions which had to be made, explicit solutions have so far been found only in special cases.

Sharing Common Pool Resources

The sharing of so-called *common-pool resources* (see, e.g., (Ostrom et al. 1994)) involves an implicit inspection problem: once agreement has been reached, the parties concerned will want to make sure that the rules are followed. The following idealized common-pool resource problem illustrates this.

The Hatfields and the McCoy's (who else!) farm neighboring areas of irrigated land, the Hatfields having first access to the water. They have agreed on a fair sharing procedure, but there exists a certain amount of mutual distrust. (In fact, the two clans have not been on speaking terms for 40 years.) Broadly speaking, farmer Hatfield has two strategies: to take more water than his fair share, with the twofold incentive of better crops and "doing-in" the McCoy's, or to stand by the agreement. Farmer McCoy, who waits his turn for the water, has the option of controlling Hatfield at some penalty (danger to life and limb) or of simply doing nothing and taking what he gets. The strategic situation is illustrated in Fig. 3.

We may normalize the payoffs to zero for legal behavior and no control. Relative to this normalization, let e be the cost to McCoy of monitoring



Inspection Games, Fig. 3 The games irrigators play (Reprinted from Avenhaus and Canty (1996) with permission of Cambridge University Press)

Hatfield's activities and a the loss he entails by not getting his fair share, but having the satisfaction of catching Hatfield out, $e > 0$, $a > 0$. Thus if Hatfield violates and McCoy monitors, McCoy's payoff is $-a$ relative to the normalization, while if he monitors but there is no violation his payoff is $-e$. Suppose, however, that McCoy's highest priority is to keep Hatfield honest. Then, necessarily, $a > e$. The worst outcome for McCoy is certainly undetected violation, with payoff $-c$, and so $c > a$. Hatfield's payoffs are simply $+d$ for undetected violation, and $-b$ for detected violation, $(b, d) > (0, 0)$. The game is depicted in bimatrix form in Fig. 4.

Solution

Just as in the passenger ticket control game, the preference arrows are cyclic. There is again a unique equilibrium in mixed strategies:

$$\beta^* = \frac{b}{b+d}, \quad t^* = \frac{e}{e+c-a}, \quad (28)$$

with payoffs to McCoy and Hatfield, respectively, given by

$$I_M(\beta^*, t^*) = -e \cdot \frac{c}{e+c-a}, \quad I_H(\beta^*, t^*) = 0. \quad (29)$$

Hatfield violates his commitment with probability $t^* > 0$ even though his payoff is the same as for behaving legally. From a moralist's viewpoint this may not be particularly satisfactory, but Hatfield's equilibrium behavior is, given the circumstances, rational.

However it was postulated that McCoy's highest priority was to keep Hatfield honest, and there is in fact a way for him to do it. Suppose that McCoy takes the initiative and *announces credibly with what precise probability* he intends to monitor Hatfield's activities. This so-called *leadership game* can no longer be expressed as a bimatrix as in Fig. 4, because McCoy's set of

pure strategies, that is, his choices of which monitoring probability $1 - \beta$ he will announce in advance, is infinite. Hatfield's set of strategies on the other hand consists of all functions which assign to each value of $1 - \beta$ the decision "take share" or "take more." The appropriate representation is the extensive form game shown in Fig. 5.

Due to its structure (formally, an extensive form game with *perfect information* in which all information sets are singletons) the game can be solved by *backward induction*. If this procedure leads to an equilibrium, then that equilibrium is said to be *subgame perfect*. A subgame perfect Nash equilibrium is one in which every subgame is also a Nash equilibrium.

The first step is to replace the outcome payoffs in Fig. 5 by their expected values, as shown in Fig. 6. The argument then proceeds as follows: Hatfield, knowing the probability $1 - \beta$ of being controlled, decides

takeshare if $0 > -b + (b + d)\beta$
indifferent if $0 = -b + (b + d)\beta$
takemore if $0 < -b + (b + d)\beta$,

(30)


since this strategy will always maximize his expected payoff. McCoy's equilibrium strategy will be shown below to be


		2	
		take share $1 - t$	take more t
1	no control β	0 0	$+d$ $-c$
	control $1 - \beta$	0 $-e$	$-b$ $-a$

Inspection Games, Fig. 4 Normal or bimatrix form of the irrigation game. *Left lower numbers* are the payoffs to player 1 (McCoy), *upper right numbers* the payoffs to player 2 (Hatfield), whereby $C > a > e$ and $(a, b, c, d, e) > (0, 0, 0, 0, 0)$. The *horizontal arrows* are incentive directions for player 2, the *vertical arrows* for player 1. The variables t and β define mixed strategies

Inspection Games,

Fig. 5 Extensive form of the inspector leadership irrigation game. McCoy begins by choosing his monitoring probability $1 - \beta$ and announces it to Hatfield, the latter then deciding whether or not to take more than his share of water. Finally, chance decides if the monitoring actually takes place. The payoffs for each possible outcome are shown at the corresponding leaf nodes, McCoy *above*, Hatfield *below*



Inspection Games,**Fig. 6** Reduced extensive form of the inspector leadership irrigation game of Fig. 5

$$\beta^* = \frac{b}{b+d}, \quad (31) \quad t^*(\beta) = \begin{cases} 0, & \text{take share} & \text{for } \beta \leq \beta^* \\ 1, & \text{take more} & \text{for } \beta > \beta^* \end{cases} \quad (32)$$

so (30) is equivalent to Hatfield's following decision:

$$\begin{array}{ll} \text{takeshare} & \text{if } \beta < \beta^* \\ \text{indifferent} & \text{if } \beta = \beta^* \\ \text{takemore} & \text{if } \beta > \beta^* \end{array} \quad (32)$$

What is Hatfield's equilibrium strategy? In order to determine it, we must first define his set of strategies a little more carefully. A typical element of the set will be a recipe which tells him, for every conceivable announcement by McCoy of a value of the non-monitoring probability β , whether or not to take more than his share. Recipe (30) is certainly one such, although it leaves an ambiguity for $\beta = \beta^*$. In that case Hatfield may decide to make his choice randomly, in other words to use a mixed strategy. In general, such a mixed strategy is given by a probability $t(\beta)$ for "take more," and the complementary probability $1 - t(\beta)$ for "take share," for all $\beta \in [0, 1]$ Hatfield's complete strategy set is therefore the set of all functions which map the unit interval onto itself, $\{t(\beta) \mid t: [0, 1] \mapsto [0, 1]\}$.

We now assert that Hatfield's equilibrium strategy t^* is in fact always a pure strategy, namely

where β^* is given by (31). This is just the conclusion we reached in Eq. (32) by backward induction, except for the case $\beta = \beta^*$. We still have to show that $t^*(\beta^*) = 0$, that is, that Hatfield stays honest at equilibrium. To do this, we now have to consider McCoy's payoffs. The expected payoff to McCoy, as a function of β , is given by


$$I_M(\beta, t^*) = \begin{cases} -e(1-\beta) & \text{if } \beta < \beta^* \\ -e(1-\beta^*)(1-t(\beta^*)) \\ +(-c\beta^* - a(1-\beta^*))t(\beta^*) & \text{if } \beta = \beta^* \\ -c\beta - a(1-\beta) & \text{if } \beta > \beta^* \end{cases} \quad (34)$$

It is plotted in Fig. 7. McCoy also wishes to maximize his expected payoff for $\beta < \beta^*$, when Hatfield stays honest, it is at least $-e$ and increases with increasing β . For $\beta > \beta^*$, when Hatfield takes more than his share, it is between $-c$ and $-a$ and certainly worse for McCoy. For $\beta = \beta^*$, McCoy's payoff is something intermediate, depending on Hatfield's behavior. The argument seems to be getting circular, but we are almost done.

McCoy's equilibrium strategy, if he has one at all, has to be β^* as given by (31): for $\beta > \beta^*$ McCoy would do better by choosing alternatively

Inspection Games,

Fig. 7 McCoy's payoff as a function of β according to (34). The • indicates the equilibrium payoff of the simultaneous game, Eq. (28)



a small β to make Hatfield act honestly, and for $\beta < \beta^*$ he could always do a little better by choosing a larger β , closer to β^* . So the only maximum of his payoff curve, as seen in Fig. 7, is at $\beta = \beta^*$. However, as Eq. (34) shows, the maximum exists only if Hatfield's equilibrium strategy is such that $t^*(\beta^*) = 0$. The unique subgame perfect equilibrium strategies must therefore be (β^*, t^*) given by (31) and (33), with payoffs

$$I_M(\beta^*, t^*) = -c \cdot \frac{d}{b+d}, \quad I_H(\beta^*, t^*) = 0. \quad (35)$$

Thus McCoy's equilibrium strategy is the same as before, as is Hatfield's payoff, the decisive difference being that Hatfield does not take more than his share of the water.

Remarks

The leadership concept was first introduced by von Stackelberg (1934) in the context of economic theory, well before game theory became a scientific discipline. In game-theoretic terminology Schelling (1960) probably was the first to formulate its importance:

A strategic move is one that influences the other person's choice in a manner favorable to one's self, by affecting the other person's expectations on how one's self will behave. One constrains the partner's choice by constraining one's own behavior. The object is to set up for one's self and communicate persuasively to the other player a mode of behavior (including conditional responses to the other's behavior) that leaves the other a simple maximization problem whose solution for him is the optimum

for one's self, and to destroy the other's ability to do the same.

Simaan and Cruz (1973) and Wölling (2002) refined the concept and formulated conditions for the existence of Nash equilibria in leadership games. Maschler was the first to apply the leadership concept to inspection games (Maschler 1966, 1967). Later on it was widely used in the analysis of IAEA verification procedures, in particular for variable sampling inspection problems, see Avenhaus et al. (1991).

The importance of deterring the inspectee from illegal behavior, or more positively, of inducing him to behave legally, depends on the specific nature of the problem. For example in the ticket control problem of Subsection "Passenger Ticket Control," maximization of intake – fares plus fines – is no doubt the highest priority of the transit authority, even though the inspector leadership concept would work here as well, at least in theory. In principal agent models the situation is similar. In the context of arms control and disarmament, on the other hand, deterrence is fundamental: the community of States party to such an agreement have a vital interest that all members adhere to its provisions. There exists a large literature on the subject; a comprehensive survey of the leadership concept in game theory in general has been given by Wölling (2002).

Future Directions

We hope that, with our chosen examples, we have given a representative, although certainly not exhaustive, overview of the concepts and models making up

the area of inspection games. At the same time we have tried to give some idea of its wide range of applications.

Of course we cannot predict future developments in the field. To be sure, there are still unsolved mathematical problems associated with present inspection models, in particular in arms control and disarmament. For example in Subsection “[Diversion of Nuclear Material](#)” it was pointed out that near real time accountancy poses fundamental difficulties that have not yet been solved satisfactorily. Active research is proceeding and interesting results may be expected.

As mentioned at the outset, in the area of environmental control the number of published investigations is surprisingly small. With the growing awareness of the importance of international agreement on environmental protection the need for effective and efficient control mechanisms will become more and more apparent. Here we expect that the inspection game approach, especially as a means of structuring verification systems, can and will play a useful role. As Barry O’Neill concludes his examination of game theory in peace and war (O’Neill 1994):

... game theory clarifies international problems exactly because they are more complicated. [...] The contribution of game models is to sort out concepts and figure out what the game might be.

Bibliography

- Avenhaus R (1997) Entscheidungstheoretische Analyse der Fahrgast-Kontrollen. *Der Nahverkehr* 9:27
- Avenhaus R, Canty MJ (1989) Re-examination of the IAEA formula for stratified attribute sampling. In: Proceedings of the 11th ESARDA symposium, JRC, Ispra, pp 351–356
- Avenhaus R, Canty MJ (1996) Compliance quantified. Cambridge University Press, Cambridge
- Avenhaus R, Canty MJ (2005) Playing for time: a sequential inspection game. *Eur J Oper Res* 167(2):474–492
- Avenhaus R, Jaech JL (1981) On subdividing material balances in time and/or space. *J Inst Nucl Manag* IV (3):24–33
- Avenhaus R, Okada A, Zamir S (1991) Inspector leadership with incomplete information. In: Selten R (ed) *Game equilibrium models*, vol IV. Springer, Heidelberg, pp 319–361
- Avenhaus R, Canty MJ, Kilgour DM, von Stengel B, Zamir S (1996) Inspection games in arms control. *Eur J Oper Res* 90:383–394
- Avenhaus R, von Stengel B, Zamir S (2002) Inspection games. In: Aumann R, Hart S (eds) *Handbook of game theory*. Elsevier, Amsterdam, pp 1947–1987
- Baiman S (1982) Agency research in managerial accounting: a survey. *J Account Lit* 1:154–213
- Baston VJ, Bostock FA (1991) A remark on the customs smuggler game. *Nav Res Logist* 41:287–293
- Bierlein D (1968) Direkte Überwachungssysteme. *Oper Res Verfahr* 6:57–68
- Bierlein D (1969) Auf Bilanzen und Inventuren basierenden Safeguards-Systeme. *Oper Res Verfahr* 6:36–43
- Borch K (1990) *Economics of insurance*. North-Holland, Amsterdam
- Cavasoglu H, Raghunathan S (2004) Configuration of detection software: a comparison of decision and game theory. *Decis Anal* 1:131–148
- Cook J, Nadeau L, Thomas LC (1997) Does cooperation in auditing matter? A comparison of a non-cooperative and a cooperative game model of auditing. *Eur J Oper Res* 103:470–482
- Derman C (1961) On minimax surveillance schedules. *Nav Res Logist* 8:415–419
- Diamond H (1982) Minimax policies for unobservable inspections. *Math Oper Res* 7(1):139–153
- Dresher M (1962) A sampling inspection problem in arms control agreements: a game theoretical analysis. Memorandum RM-2972-ARPA. RAND Corporation, Santa Monica
- Dye RA (1986) Optimal monitoring policies in agencies. *RAND J Econ* 17:339–350
- Ferguson TS, Melolidakis C (1998) On the inspection game. *Nav Res Logist* 45:327–334
- Garnaev AY (1991) A generalized inspection game. *Nav Res Logist* 28:171–188
- Goutal P, Garnaev A, Garnaeva G (1997) On the infiltration game. *Int J Game Theory* 26(2):215–221
- Höpfinger E (1971) A game-theoretic analysis of an inspection problem, University of Karlsruhe (unpublished manuscript)
- Höpfinger E (1974) Zuverlässige Inspektionsstrategien. *Z Wahrscheinlichkeitstheorie Verw Geb* 31:35–46
- Hozaki R, Kuhdoh D, Komiyama T (2006) An inspection game: taking account of fulfillment probabilities of players. *Nav Res Logist* 53:761–771
- IAEA (1972) The structure and content of agreements between the agency and states required in connection with the treaty on the non-proliferation of nuclear weapons. IAEA, Vienna, INF/CIRC 153 (corrected)
- IAEA (1997) Model protocol additional to the agreement (s) between state(s) and the international atomic energy agency for the application of safeguards. IAEA, Vienna, INF/CIRC 140
- Kanodia CS (1985) Stochastic and moral hazard. *J Account Res* 23:175–293

- Kilgour DM (1992) Site selection for on-site inspection in arms control. *Arms Control* 13:439–462
- Krieger T (2008) On the asymptotic behavior of a discrete time inspection game. *Math Model Anal* 13(1):37–46
- Kuhn HW (1953) Extensive games and the problem of information. In: Kuhn HW, Tucker AW (eds) *Contributions to the theory of games*, vol II. Princeton University Press, Princeton, pp 193–216
- Maschler M (1966) A price leadership method for solving the inspector's non-constant-sum game. *Nav Res Logist* 13:11–33
- Maschler M (1967) The inspector's non-constant-sum-game: its dependence on a system of detectors. *Nav Res Logist* 14:275–290
- Morris P (1994) *Introduction to game theory*. Springer, New York
- Nash JF (1951) Non-cooperative games. *Ann Math* 54:286–295
- O'Neill B (1994) Game theory models of peace and war. In: Aumann R, Hart S (eds) *Handbook of game theory*. Elsevier, Amsterdam, pp 995–1053
- Ostrom E, Gardner R, Walker J (1994) *Rules, games and common pool resources*. University of Michigan Press, Ann Arbor
- Owen G (1968) *Game theory*. W. B. Saunders, Philadelphia
- Pavlovic L (2002) More on the search for an infiltrator. *Nav Res Logist* 49:1–14
- Rinderle K (1996) Mehrstufige sequentielle Inspektionsspiele mit statistischen Fehlern erster und zweiter Art. Kovac, Hamburg
- Rohatgi VK (1976) *An introduction to probability theory and mathematical statistics*. Wiley, New York
- Rothenstein D, Zamir S (2002) Imperfect inspection games over time. *Ann Oper Res* 109:175–192
- Sakaguchi M (1994) A sequential game of multi-opportunity infiltration. *Math Jpn* 39:157–166
- Schelling TC (1960) *The strategy of conflict*. Harvard University Press, Cambridge, MA
- Simaan M, Cruz JB (1973) On the Stackelberg strategy in nonzero-sum games. *J Optim Theory Appl* 11(5): 533–555
- von Neumann J, Morgenstern O (1947) *Theory of games and economic behavior*. Princeton University Press, Princeton
- von Stackelberg H (1934) *Marktform und Gleichgewicht*. Springer, Vienna
- von Stengel B (1991) Recursive inspection games, Report No. S 9106. Computer Science Faculty, Armed Forces University Munich
- Stewart KB (1971) A cost-effectiveness approach to inventory verification. In: *Proceedings of the IAEA symposium on safeguards techniques*, vol II. International Atomic Energy Agency, Vienna, pp 387–409
- Thomas MU, Nisgav Y (1976) An infiltration game with time-dependent payoff. *Nav Res Logist* 23:297–320
- Wilks TJ, Zimbelman MF (2004) Using game theory and strategic reasoning concepts to prevent and detect fraud. *Account Horiz* 18(3):173–184
- Wölling A (2002) *Das Führerschaftsprinzip bei Inspektionsspielen*. Kovac, Hamburg

Principal-Agent Models

Inés Macho-Stadler¹ and David Pérez-Castrillo²

¹Universitat Autònoma de Barcelona and
Barcelona GSE, Barcelona, Spain

²Dept. of Economics and Economic History,
Universitat Autònoma de Barcelona and
Barcelona GSE, Barcelona, Spain

Article Outline

Glossary
Definition of the Subject
Introduction
The Base Game
Moral Hazard
Adverse Selection
Future Directions
Bibliography

Keywords

Incentives · Contracts · Asymmetric
Information · Moral Hazard · Adverse
Selection

Glossary

Adverse Selection (Hidden Information) The term adverse selection was originally used in insurance. It describes a situation where, as a result of private information, the insured are more likely to suffer a loss than the uninsured (such as offering a life insurance contract at a given premium may imply that only the people with a risk of dying over the average take it). In principal-agent models, we say that there is an adverse selection problem when the ignorant party lacks information while negotiating a contract, in such a way that the asymmetry is *previous* to the relationship.

Asymmetric Information In a relationship or a transaction, there is asymmetric information when one party has more or better information than the other party concerning relevant characteristics of the relationship or the transaction. There are two types of asymmetric information problems: moral hazard and adverse selection.

Information Economics Information economics studies how information and its distributions among the players affect economic decisions.

Moral Hazard (Hidden Action) The term moral hazard initially referred to the possibility that the redistribution of risk (such as insurance which transfers risk from the insured to the insurer) changes people's behavior. This term, which has been used in the insurance industry for many years, was studied first by Kenneth Arrow. In principal-agent models, the term moral hazard is used to refer to all environments where the ignorant party lacks information about the behavior of the other party *once* the agreement has been signed, in such a way that the asymmetry arises *after* the contract is settled.

Principal Agent The principal-agent model identifies the difficulties that arise in situations where there is asymmetric information between two parties and finds the best contract in such environments. The “principal” is the name used for the contractor, while the “agent” corresponds to the contractee. Both principal and agent could be individuals, institutions, organizations, or decision centers. The optimal solutions propose mechanisms that try to align the interests of the agent with those of the principal, such as piece rates or profit sharing, or that induce the agent to reveal the information, such as self-reporting contracts.

Definition of the Subject

Principal-agent models provide the theory of *contracts under asymmetric information*. Such a theory analyzes the characteristics of optimal contracts and the variables that influence these characteristics, according to the behavior and information of the parties to the contract. This approach

has a close relation to *game theory* and *mechanism design*: it analyzes the strategic behavior by agents who hold private information and proposes mechanisms that minimize the inefficiencies due to such strategic behavior. The costs incurred by the principal (the contractor) to ensure that the agents (the contractees) will act in her interest are some type of *transaction cost*. These costs include the tasks of investigating and selecting appropriate agents, gaining information to set performance standards, monitoring agents, bonding payments by the agents, and residual losses.

Principal-agent theory (and *information economics* in general) is possibly the area of economics that has evolved the most over the past 25 years. It was initially developed in parallel with the new economics of *industrial organization*, although its applications include now almost all areas in economics, from finance and political economy to growth theory.

Some early papers centered on incomplete information in insurance contracts, and more particularly on moral hazard problems, are Spence and Zeckhauser (1971) and Ross (1973). The theory soon generalized to dilemmas associated with contracts in other contexts (Harris and Raviv 1978; Jensen and Meckling 1976). It was further developed in the mid-1970s by authors such as Pauly (1968, 1974), Mirrlees (1975), Harris and Raviv (1979), and Holmström (1979). Arrow (1985) worked on the analysis of the optimal incentive contract when the agent's effort is not verifiable.

A particular case of adverse selection is the one where the type of the agent relates to his valuation of a good. Asymmetric information about buyers' valuation of the objects sold is the fundamental reason behind the use of auctions. Vickrey (1961) provides the first formal analysis of the first- and second-prize auctions. Akerlof (1970) highlighted the issue of adverse selection in his analysis of the market for secondhand goods. Further analyses include the early work of Mirrlees (1971), Spence (1974), Rothschild and Stiglitz (1976), Mussa and Rosen (1978), Baron and Myerson (1982), and Guesnerie and Laffont (1984).

The importance of the topic has also been recognized by the Nobel Foundation. James A. Mirrlees and William Vickrey were awarded with the Nobel

Prize in Economics in 1996 "for their fundamental contributions to the economic theory of incentives under asymmetric information." Five years later, in 2001, George A. Akerlof, A. Michael Spence, and Joseph E. Stiglitz also obtained the Nobel Prize in Economics "for their analyses of markets with asymmetric information."

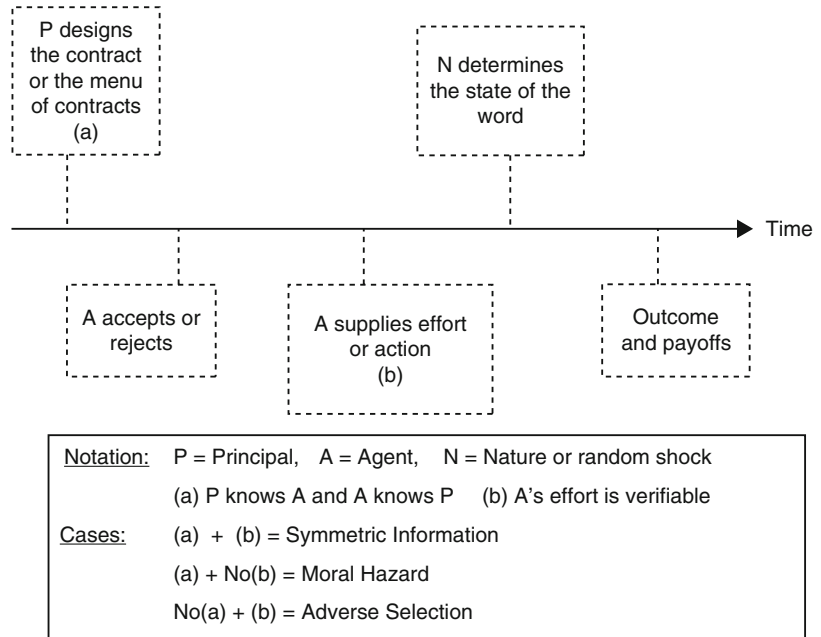
Introduction

The objective of the principal-agent literature is to analyze situations in which a contract is signed under asymmetric information, that is, when one party knows certain relevant things of which the other party is ignorant. The simplest situation concerns a bilateral relationship: the contract between one principal and one agent. The objective of the contract is for the agent to carry out actions on behalf of the principal and to specify the payments that the principal will pass on to the agent for such actions.

In the literature, it is always assumed that the principal is in charge of designing the contract. The agent receives an offer and decides whether or not to sign the contract. He will accept it whenever the utility obtained from it is greater than the utility that the agent would get from not signing. This utility level that represents the agent's outside opportunities is his *reservation utility*. In order to simplify the analysis, it is assumed that the agent cannot make a counteroffer to the principal. This way of modeling implicitly assumes that the principal has all the bargaining power, except for the fact that the reservation utility can be high in those cases where the agent has excellent outside opportunities.

If the agent decides not to sign the contract, the relationship does not take place. If he does accept the offer, then the contract is implemented. It is crucial to notice that the contract is a reliable promise by both parties, stating the principal and agent's obligations for all (contractual) contingencies. It can only be based on *verifiable variables*, that is, those for which it is possible for a third party (a court) to verify whether the contract has been fulfilled. When some players know more than others about relevant variables, we have a

Principal-Agent Models, Fig. 1 The figure summarizes the timing of the relationship and the three cases as a function of the information available to the participants



situation with asymmetric information. In this case, *incentives* play an important role (Fig. 1).

Given the description of the game played between the principal and agent, we can summarize its timing in the following steps:

1. The principal designs the contract (or set of contracts) that she will offer to the agent, the terms of which are not subject to bargaining.
2. The alternatives opened to the agent are to accept or to reject the contract. The agent accepts it if he desires so, that is, if the contract guarantees him greater expected utility than any other (outside) opportunities available to him.
3. The agent carries out an action or effort on behalf of the principal.
4. The outcome is observed and the payments are done.

From these elements, it can be seen that *the agent's objectives may be in conflict with those of the principal*. When the information is asymmetric, the informed party tries to take advantage, while the uninformed player tries to control this behavior via the contract. Since a principal-agent problem is a sequential game, the solution

concept to use is *subgame (Bayesian) perfect equilibrium*.

The setup gives rise to three possible scenarios:

1. The *symmetric information* case, where the two players share the same information, even if they both may ignore some important elements (some elements may be uncertain).
2. The *moral hazard* case, where the asymmetry of information arises once the contract has been signed: the decision or the effort of agent is not verifiable and hence it cannot be included in the contract.
3. The *adverse selection* case, where the asymmetry of information is previous to the signature of the contract: a relevant characteristic of the agent is not verifiable and hence the principal cannot include it in the contract.

To see an example of moral hazard, consider a laboratory or research center (the principal) that contracts a researcher (the agent) to work on a certain project. It is difficult for the principal to distinguish between a researcher who is thinking about how to push the project through and a researcher who is thinking about how to organize

his evening. It is precisely this difficulty in controlling effort inputs, together with the inherent uncertainty in any research project, what generates a moral hazard problem, which is a nonstandard labor market problem.

For an example of adverse selection, consider a regulator who wants to set the price of the service provided by a public monopoly equal to the average costs in the firm (to avoid subsidies). This policy (as many others) is subject to important informational requirements. It is not enough that the regulator asks the firm to reveal the required information in order to set the adequate price, since the firm would attempt to take advantage of the information. Therefore, the regulator should take this problem into account.

The Base Game

Consider a contractual relationship between a principal and an agent, who is contracted to carry out a task. The relationship allows a certain *result* to be obtained, whose monetary value will be referred to as x . For the sake of exposition, the set of possible results X is assumed to be finite, $X = \{x_1, \dots, x_n\}$. The final result depends on the *effort* that the agent devotes to the task, which will be denoted by e , and the value of a random variable for which both participants have the same prior distribution. The probability of result x_i conditional on effort e can be written as

$$\text{Prob}[x = x_i|e] = p_i(e), \quad \text{for } i \in \{1, 2, \dots, n\},$$

with $\sum_{i=1}^n p_i(e) = 1$. Let us assume that $p_i(e) > 0$ for all e, i , which implies that no result can be ruled out for any given effort level.

The base game is the reference situation, where the principal and agent have the same information (even the one concerning the random component that affects the result). Since uncertainty exists, participants react to risk. Risk preferences are expressed by the shape of their *utility functions* (of the von Neumann-Morgenstern type). The

principal, who owns the result and must pay the agent, has preferences represented by the utility function

$$B(x - w),$$

where w represents the payoff made to the agent. $B(\cdot)$ is assumed to be increasing and concave: $B' > 0$, $B'' \leq 0$ (where the primes represent, respectively, the first and second derivatives). The concavity of the function $B(\cdot)$ indicates that the principal is either risk neutral or risk averse.

The agent receives a monetary payoff for his participation in the relationship, and he supplies an effort which implies some cost to him. For the sake of simplicity, we represent his utility function as

$$U(w, e) = u(w) - v(e),$$

additively separable in the components w and e . This assumption implies that the agent's risk aversion does not vary with the effort he supplies (many results can be generalized for more general utility functions). The utility derived from the wage, $u(w)$, is increasing and concave:

$$u'(w) > 0, \quad u''(w) \leq 0,$$

and thus the agent may be either risk neutral, $u''(w) = 0$, or risk averse, $u''(w) < 0$. In addition, greater effort means greater disutility. We also assume that the marginal disutility of effort is not decreasing: $v'(e) > 0$, $v''(e) \geq 0$.

A *contract* can only be based in verifiable information. In the base game, it can depend on the characteristics of the principal and the agent and includes both the effort e that the principal demands from the agent and the wages $\{w(x_i)\}_{i=1, \dots, n}$.

If the agent rejects the contract, he will have to fall back on the outside opportunities that the market offers him. These other opportunities, which by comparison determine the limit for participation in the contract, are summarized in the agent's *reservation utility*, denoted by $\underline{U}$. So the agent will accept the contract as long as he earns an expected utility equal or higher than his

reservation utility. Since the principal's problem is to design a contract that the agent will accept (by backward induction), the optimal contract must satisfy the *participation constraint*, and it is the solution to the following maximization problem:

$$\begin{aligned} \text{Max}_{[e, \{w(x_i)\}_{i=1, \dots, n}]} & \sum_{i=1}^n p_i(e) B(x_i - w(x_i)) \text{ s.t.} \\ & \sum_{i=1}^n p_i(e) u(w(x_i)) - v(e) \geq \underline{U}. \end{aligned}$$

The above problem corresponds to a *Pareto optimum* in the usual sense of the term. The solution to this problem is conditional on the value of the parameter $\underline{U}$, so that even those cases where the agent can keep a large share of the surplus are taken into account.

The principal's program is well behaved with respect to payoffs given the assumptions on $u(w)$. Hence the Kuhn-Tucker conditions will be both necessary and sufficient for the global solution of the problem. However, we cannot ascertain the concavity (or quasi-concavity) of the functions with respect to effort given the assumptions on $v(e)$, because these functions also depend on all the $p_i(e)$. Hence it is more difficult to obtain global conclusions with respect to this variable.

Let us denote by e° the efficient effort level. From the first-order Kuhn-Tucker conditions with respect to the wages in the different contingencies, we can analyze the associated payoffs $\{w^\circ(x_i)_{i=1, \dots, n}\}$. We obtain the following condition:

$$\lambda^\circ = \frac{B'(x_i - w^\circ(x_i))}{u'(w^\circ(x_i))}, \quad \text{for all } i \in \{1, 2, \dots, n\},$$

where λ° is the multiplier associated with the participation constraint. When the agent's utility is additively separable, the participation constraint binds (λ° is positive). The previous condition equates marginal rates of substitution and indicates that the optimal distribution of risk requires that the

ratio of marginal utilities of the principal and the agent to be constant irrespective of the final result.

If the principal is risk neutral ($B''(\cdot) = 0$), then the optimal contract has to be such that $u'(w^\circ(x_i)) = \text{constant}$ for all i . In addition, if the agent is risk averse ($u''(\cdot) < 0$), he receives the same wage, say w° , in all contingencies. This wage only depends on the effort demanded and is determined by the participation constraint. If the agent is risk neutral ($u''(\cdot) = 0$) and the principal is risk averse ($B''(\cdot) < 0$), then we are in the opposite situation. In this scenario, the optimal contract requires the principal's profit to be independent of the result. Consequently, the agent bears all the risk, insuring the principal against variations in the result. When both the principal and the agent are risk averse, each of them needs to accept a part of the variability of the result. The precise amount of risk that each of them supports depends on their degrees of risk aversion. Using the Arrow-Pratt measure of absolute risk aversion $r_p = -B''/B'$ and $r_a = -u''/u'$ for the principal and the agent, respectively, we can show that

$$\frac{dw^\circ}{dx_i} = \frac{r_p}{r_p + r_a},$$

which indicates how the agent's wage changes given an increase in the result x_i . Since $r_p/(r_p + r_a) \in (0, 1)$, when both participants are risk averse, the agent only receives a part of the increased result via a wage increase. The more risk averse is the agent, that is to say, the greater is r_a , the less the result influences his wage. On the other hand, as the risk aversion of the principal increases, greater r_p , changes in the result correspond to more important changes in the wage.

Moral Hazard

Basic Moral Hazard Model

Here we concentrate on *moral hazard*, which is the case in which the informational asymmetry relates to the agent's behavior during the relationship. We analyze the optimal contract when the agent's effort is not verifiable. This implies that

effort cannot be contracted upon, because in case of the breach of contract, no court of law could know if the contract had really been breached or not. There are many examples of this type of situation. A traditional example is accident insurance, where it is very difficult for the insurance company to observe how careful a client has been to avoid accidents.

The principal will state a contract based on any signals that reveal information on the agent's effort. We will assume that only the result of the effort is verifiable at the end of the period and, consequently, it will be included in the contract. However, if possible, the contract should be contingent on many other things. Any information related to the state of nature is useful, since it allows better estimations of the agent's effort thus reducing the risk inherent in the relationship. This is known as the *sufficient statistic result*, and it is perhaps the most important conclusion in the moral hazard literature (Holmström 1979). The empirical content of the sufficient statistic argument is that a contract should exploit all available information in order to filter out risk optimally.

The timing of a moral hazard game is the following. In the first place, the principal decides what contract to offer the agent. Then the agent decides whether or not to accept the relationship, according to the terms of the contract. Finally, if the contract has been accepted, the agent chooses the effort level that he most desires, given the contract that he has signed. This is a free decision by the agent since effort is not a contracted variable. Hence, the principal must bear this in mind when she designs the contract that defines the relationship.

To better understand the nature of the problem faced by the principal, consider the case of a risk-neutral principal and a risk-averse agent, which implies that, under the symmetric information, the optimal contract is to completely insure the agent. However, if the principal proposes this contract when the agent's effort is not a contracted variable, once he has signed the contract, the agent will exert the effort level that is most beneficial for him. Since the agent's wage does not depend on his effort, he will use the lowest possible effort.

The idea underlying an incentive contract is that the principal can make the agent interested in the consequences of his behavior by making his payoff dependent on the result obtained. Note that this has to be done at the cost of distorting the optimal risk sharing among both participants. The trade-off between *efficiency*, in the sense of the optimal distribution of risk, and *incentives* determines the optimal contract.

Formally, since the game has to be solved by backward induction, the *optimal contract* under moral hazard is the solution to the maximization problem:

$$\begin{aligned} & \text{Max}_{[e, \{w(x_i)\}_{i=1, \dots, n}]} \sum_{i=1}^n p_i(e) B(x_i - w(x_i)) \text{ s.t.} \\ & \sum_{i=1}^n p_i(e) u(w(x_i)) - v(e) \geq \underline{U} \\ & e \in \text{Arg Max}_e \left\{ \sum_{i=1}^n p_i(\hat{e}) u(w(x_i)) - v(\hat{e}) \right\}. \end{aligned}$$

The second restriction is the *incentive compatibility constraint* and the first restriction is the participation constraint. The incentive compatibility constraint, and not the principal as under symmetric information, determines the effort of the agent.

The first difficulty in solving this program is related to the fact that the incentive compatibility constraint is a maximization problem. The second difficulty is that the expected utility may fail to be concave in effort. Hence, to use the first-order condition of the incentive compatibility constraint may be incorrect. In spite of this, there are several ways to proceed when facing this problem. (a) Grossman and Hart (1983) propose to solve it in steps, identifying first the optimal payment mechanism for any effort and then, if possible, the optimal effort. This can be done since the problem is concave in payoffs. (b) The other possibility is to consider situations where the agent's maximization problem is well defined. One possible scenario is when the set of possible efforts is finite, in which case the incentive compatibility constraint takes the form of a finite set of inequalities. Another scenario is to write the incentive compatibility as the first-order condition of the

maximization problem and introduce assumptions that allow doing it. The last solution is known as the *first-order approach*.

Let us assume that the first-order approach is adequate, and substitute the incentive compatibility constraint in the previous program by

$$\sum_{i=1}^n p'_i(\hat{e})u(w(x_i)) - v'(\hat{e}) = 0.$$

Solving the principal program with respect to the payoff scheme and denoting by λ (resp., μ) the Lagrangian multiplier of the participation constraint (resp., the incentive compatibility constraint), we obtain that for all i :

$$\frac{1}{u'(w(x_i))} = \lambda + \mu \frac{p'_i(e)}{p_i(e)}.$$

This condition shows that the wage should not depend at all on the value that the principal places on the result. It depends on the results as a measure of how informative they are as to effort, in order to serve as an incentive for the agent. The wage will be increasing in the result as long as the result is increasing in effort. Hence, it is optimal that the wage will be increasing in the result only in particular cases. The necessary condition for a wage to be increasing with results is $p'_i(e)/p_i(e)$ to be decreasing in i . In statistics, this is called the *monotonous likelihood quotient property*. It is a strong condition; for example, first-order stochastic dominance does not guarantee the monotonous likelihood property.

Extensions of Moral Hazard Models

The basic moral hazard setup, with a principal hiring and an agent performing effort, has been extended in several directions to take into account more complex relationships.

Repeated Moral Hazard

Certain relationships in which a moral hazard problem occurs do not take place only once, but they are repeated over time (e.g., work relationships, insurance, etc.). The duration aspect (the repetition) of the relationship gives rise to new elements that are absent in static models.

Radner (1981) and Rubinstein and Yaari (1983) consider infinitely repeated relationships and show that frequent repetition of the relationship allows us to converge toward the efficient solution. Incentives are not determined by the payoff scheme contingent on the result of each period, but rather on average effort, and the information available is very precise when the number of periods is large. A sufficiently threatening punishment, applied when the principal believes that the agent on average does not fulfill his task, may be sufficient to dissuade him from shirking.

When the relationship is repeated a finite number of times, the analysis of the optimal contract concentrates on different issues relating long-term agreements and short-term contracts. Note that in a repeated setup, the agent's wage and the agent's consumption in a period need not be equal. Lambert (1983), Rogerson (1985), and Chiappori and Macho-Stadler (1990) show that long-term contracts have memory (i.e., the payoffs in any single period will depend on the results of all previous periods) since they internalize agent's consumption over time, which depends on the sequence of payments received (as a function of the past contingencies). Malcomson and Spinnewyn (1988), Fudenberg et al. (1990), and Rey and Salanié (1990) study when the optimal long-term contract can be implemented through the sequence of optimal short-term contracts. Chiappori et al. (1994) show that, in order for the sequence of optimal short-term contracts to admit the same solution as the long-term contract, two conditions must be met. First, the optimal sequence of single-period contracts should have memory. That is why, when the reservation utility is invariant (is not history dependent), the optimal sequence of short-term contracts will not replicate the long-term optimum unless there exist means of smoothing consumption, that is, the agent has access to credit markets. Second, the long-term contract must be renegotiation proof. A contract is said to be renegotiation proof if at the beginning of any intermediate period, no new contract or renegotiation that would be preferred by all participants is possible. When the long-term contract is not renegotiation proof (i.e., if it is not possible for participants to change the clauses of the contract at a certain

moment of time even if they agree), it cannot coincide with the sequence of short-term contracts.

One Principal and Several Agents

When a principal contracts with more than one agent, the stage where agents exert their effort, which is translated into the incentive compatibility, depends on the *game among the agents*. If the agents behave as a coordinated and cooperating group, then the problem is similar to the previous one where the principal hires a team. A more interesting case appears when agents play a noncooperative game and their strategies form a Nash equilibrium.

Holmström (1979) and Mookherjee (1984), in models where there is *personalized information* about the output of each agent, show that the principal is interested in paying each agent according to his own production and that of the other agents if these other results can inform on the actions of the agent at hand. Only if the results of the other agents do not add information or, in other words, if an agent's result is a *sufficient statistic* for his effort, then he will be paid according to his own result.

When the only verifiable outcome is the final result of teamwork (*joint production models*), the optimal contract can only depend on this information, and the conclusions are similar to those obtained in models with only one agent. Alchian and Demsetz (1972) and Holmström (1982) show that joint production cannot lead to efficiency when all the income is distributed among the agents, i.e., if the budget constraint always binds. Another player should be contracted to control the productive agents and act as the residual claimant of the relationship.

Tirole (1986) and Laffont (1990) have studied the effect of coalitions among the agents in an organization on their payment scheme. If collusion is bad for the organization, it adds another dimension of moral hazard (the colluding behavior). The principal may be obliged to apply rules that are collusion proof, which implies more constraints and simpler contracts (more bureaucratic). When coordination can improve the input of a group of agents, the optimal contract has to find payment methods that strengthen group work (see

Itoh 1990; Macho-Stadler and Pérez-Castrillo 1993).

Another principal's decision when she hires several agents is the organization with which she will relate. This includes such fundamental decisions as how many agents to contract and how should they be structured. These issues have been studied by Demski and Sappington (1986), Melumad and Reichelstein (1987), and Macho-Stadler and Pérez-Castrillo (1998).

Holmström and Milgrom (1991) analyze a situation in which the agent carries out *several tasks*, each one of which gives rise to a different result. They study the optimal contract when tasks are complementary (in the sense that exerting effort in one reduces the costs of the other) or substitutes. Their model allows to build a theory of job design and to explain the relationship among responsibility and authority.

Several Principals and One Agent

When one agent works for (or signs his contracts with) several principals simultaneously (*common agency* situation), in general, the principals are better off if they cooperate. When the principals are not able to achieve the coordination and commitment necessary to act as a single individual and they do not value the results in the same way, they each demand different efforts or actions from the agent. Bernheim and Whinston (1986) show that the effort that principals obtain when they do not cooperate is less than the effort that would maximize their collective profits. However, the final contract that is offered to the agent minimizes the cost of getting the agent to choose the contractual effort.

Adverse Selection

Basic Adverse Selection Model

Adverse selection is the term used to refer to problems of asymmetric information that appear before the contract is signed. The classic example of Akerlof (1970) illustrates very well the issue: the buyer of a used car has much less information about the state of the vehicle than the seller. Similarly, the buyer of a product knows how much he

appreciates the quality, while the seller only has statistical information about a typical buyer's taste (Mussa and Rosen 1978), or the regulated firm has more accurate information about the marginal cost of production than the regulator.

A convenient way to model adverse selection problems is to consider that the agent can be of different *types* and that the agent knows his type before any contract is signed while the principal does not know it. In the previous examples, the agent's type is the particular quality of the used car, the level of appreciation of quality, or the firm's marginal cost. How can the principal deal with this informational problem? Instead of offering just one contract for every (or several) type of agents, she can propose several contracts so that each type of agent chooses the one that is best for him. A useful result in this literature is the *revelation principle* (Gibbard 1973; Green and Laffont 1977; Myerson 1979) that states that any mechanism that the principal can design is equivalent to a *direct revelation mechanism* by which the agent is asked to reveal his type and a contract is offered according to his declaration. That is, a direct revelation mechanism offers a *menu of contracts* to the agent (one contract for each possible type), and the agent can choose any of the proposed contracts. Clearly, the mechanism must give the agent the right incentives to choose the appropriate contract, that is, it must be a *self-selection mechanism*. Menus of contracts are not unusual. For instance, insurance companies offer several possible insurance contracts between which clients may freely choose their most preferred. For example, car insurance contracts can be with or without deductible clauses. The second goes to more risk averse or more frequent drivers, while deductibles attract less risk averse or less frequent drivers.

Therefore, the timing of an adverse selection game is the following. In the first place, the agent's characteristics (his "type") are realized, and only the agent learns them. Then, the principal decides the menu of contracts to offer to the agent. Having received the proposal, the agent decides which one of the contracts (if any) to accept. Finally, if one contract has been accepted, the agent chooses the predetermined effort and receives the corresponding payment.

A simple model of adverse selection is the following. A risk-neutral principal contracts an agent (who could be risk neutral or risk averse) to carry out some verifiable effort on her behalf. Effort e provides an expected payment to the principal of $\Pi(e)$, with $\Pi'(e) > 0$ and $\Pi''(e) < 0$. The agent could be either of *two types* that differ with respect to the disutility of effort, which is $v(e)$ for type G(good) and $kv(e)$, with $k > 1$, for type B (bad). Hence, the agent's utility function is either $U^G(w, e) = u(w) - v(e)$ or $U^B(w, e) = u(w) - kv(e)$. The principal considers that the probability for an agent to be type G is q , where $0 < q < 1$.

The principal designs a menu of contracts $\{(e^G, w^G), (e^B, w^B)\}$, where (e^G, w^G) is directed toward the most efficient type of agent, while (e^B, w^B) is intended for the least efficient type. For the menu of contracts to be a sensible proposal, the agent must be better off by truthfully revealing his type than by deceiving the principal. The principal's problem is therefore to maximize her expected profits subject to the restrictions that, (a) after considering the contracts offered, the agent decides to sign with the principal (*participation constraints*), and (b) each agent chooses the contract designed for his particular type (*incentive compatibility constraints*):

$$\begin{aligned} & \text{Max}_{[(e^G, w^G), (e^B, w^B)]} q[\Pi(e^G) - w^G] + (1 - q) \\ & \quad \times [\Pi(e^B) - w^B] \text{ s.t. } u(w^G) - v(e^G) \\ & \geq \underline{u}(w^B) - kv(e^B) \geq \underline{u}(w^G) - v(e^G) \\ & \geq u(w^B) - v(e^B) \quad u(w^B) - kv(e^B) \\ & \geq u(w^G) - kv(e^G). \end{aligned}$$

The main characteristics of the optimal contract menu $\{(e^G, w^G), (e^B, w^B)\}$ are the following:

1. The contract offered to the good agent (e^G, w^G) is efficient (*non-distortion at the top*). The optimal salary w^G however is higher than under symmetric information: this type of agent receives an *informational rent*. That is, the most efficient agent profits from his private information, and in order to reveal this

information, he has to receive a utility greater than his reservation level.

2. The participation condition binds for the agent when he has the highest costs (he just receives his reservation utility). Moreover, a distortion is introduced into the efficiency condition for this type of agent. By distorting, the principal loses efficiency with respect to type-*B* agents, but she pays less informational rent to the *G* types.

Principals Competing for Agents in Adverse Selection Frameworks

Starting with the pioneer work by Rothschild and Stiglitz (1976) on insurance markets, there have been many studies on markets with adverse selection problems where there is competition among principals to attract agents. We move from a model where one principal maximizes her profits subject to the above constraints, to a game theory environment where each principal has to take into account the actions by others when deciding which contract to offer. In this case, the adverse selection problem may be so severe that we may find ourselves in situations in which no equilibrium exists.

To highlight the main results in this type of models, consider a simple case in which there are two possible risk-averse agent types: good (*G*) and bad (*B*) with *G* being more productive than *B*. In particular, we assume that *G* is more careful than *B*, in the sense that he commits fewer errors. When the agent exerts effort, the result could be either a success (*S*) or a failure (*F*). The probability that it is successful is p^G when the agent is type *G* and p^B when he is type *B*, where $p^G > p^B$. The principal values a successful result more than a failure. The result is observable, so that the principal can pay the agent according to the result, if she so desires.

There are several risk-neutral principals. Therefore, we look for the set of equilibrium contracts in the game played by principals competing to attract agents. Equilibrium contracts must satisfy that there does not exist a principal who can offer a different contract that would be

preferred by all or some of the agents and that gives that principal greater expected profits. This is why, if information was symmetric, the equilibrium contracts would be characterized by the following properties: (i) principals' expected profits are zero, and (ii) each contract must be efficient. Hence the agent receives a fixed contract insuring him against random events. In particular, the equilibrium salary that the agent receives under symmetric information is higher when he is of type *G* than when he is of type *B*.

When the principals cannot observe the type of the agent, the previous contracts can no longer be an equilibrium: all the agents would claim to be a good type. An equilibrium contract pair $\{C^G, C^B\}$ must satisfy the condition that no principal can add a new contract that would give positive expected profits to the agents that prefer this new contract to C^G and C^B . If the equilibrium contracts for the two agent types turn out to be the same, that is, there is only one contract that is accepted by both agent types, then the equilibrium is said to be *pooling*. On the other hand, when there is a different equilibrium contract for each agent type, then we have a *separating* equilibrium. In fact, pooling equilibria never exist, since pooling contracts always give room for a principal to propose a profitable contract that would only be accepted by the *G* types (the best agents). If an equilibrium does exist, it must be such that each type of agent is offered a different contract.

If the probability that the agent is "good" is large enough, then a separating equilibrium does not exist either. That is, an adverse selection problem in a market may provoke the absence of any equilibrium in that market. When separating equilibria do exist, the results are similar to the ones under moral hazard in spite of the differences in the type of asymmetric information and in the method of solving. That is, contingent payoffs are offered to the best agent to allow the principal to separate them from the less efficient ones. In this equilibrium, the least efficient agents obtain the same expected utility (and even sign the same contract) as under symmetric information, while the best agents lose expected utility due to the asymmetric information.

Extensions of Adverse Selection Models

Repeated Adverse Selection

In this extension, we consider whether the repetition of the relationship during several periods helps the principal and how it influences the form of the optimal contract. Note first that if the agent's private information is different in each period and the information is not correlated among periods, then any current information revealed does not affect the future, and hence the repeated problem is equivalent to simple repetition of the initial relationship. The optimal intertemporal contract will be the sequence of optimal single-period contracts.

Consider the opposite situation where the agent's type is constant over time. If the agent decides to reveal his type truthfully in the first period, then the principal is able to design efficient contracts that extract all surpluses from the agent. Hence, the agent will have very strong incentives to misrepresent his information in the early periods of the relationships. In fact, Baron and Besanko (1984) show that if the principal can commit herself with a contract that covers all the periods, then the optimal contract is the repetition of the optimal static contract. This implies that the contract is not *sequentially rational*, and it is also *non-robust to renegotiation*: once the first period is finished, the principal "knows" the agent's type and a better contract for both parties is possible.

It is often the case that the principal cannot commit not to renegotiate a long-term contract. Laffont and Tirole (1988) show that, in this case, it may be impossible to propose perfect revelation (separating) contracts in the first periods. This is known as the *ratchet effect*. Also, Freixas et al. (1985) and Laffont and Tirole (1987) have proven that, even when separating contracts exist, they may be so costly that they are often nonoptimal, and we should expect that information be revealed progressively over time. Baron and Besanko (1987) and Laffont and Tirole (1990) also introduce frameworks in which it is possible to propose perfect revelation contracts, but they are not optimal.

Relationships with Several Agents: Auctions

One particularly interesting case of relationship among one principal and several agents is that of a seller who intends to sell one or several items to several interested buyers, where buyers have private information about their valuation for the item (s). A very popular selling mechanism in such a case is an *auction*. As Klemperer (2004) would put it, auction theory is one of economics' success stories in both practical and theoretical terms. Art galleries generally use *English auctions* where the agents bid "upwards," while fish markets are generally examples of *Dutch auctions* where the seller reduces the price of the good until someone stops the auction by buying. Public contracts are generally awarded through (*first price* or *second price*) *sealed-bid auctions* where buyers introduce their bid in a closed envelope, the good is sold to the highest bidder, and the prize is either the winner's own bid or the second highest bid.

Vickrey (1961, 1962) was the first to establish the key result in auction theory, the *revenue equivalence theorem* which, subject to some reasonable conditions, says that the seller can expect equal profits on average from all the above (and many other) types of auctions and that buyers are also indifferent among them all. Auctions are efficient, since the buyer who ends up with the object is the one with the highest valuation. Hence, the existence of private information does not generate any distortions with respect to who ends up getting the good, but the revenue of the seller is less than under symmetric information.

Myerson (1981) solves the general mechanism design problem of a seller who wants to maximize her expected revenue, when the bidders have independent types and all agents are risk neutral. In general, the optimal auction is more complex than the traditional English (second price) or Dutch (first price) auction. His work has been extended by many other authors. When the buyers' types are affiliated (i.e., they are not negatively correlated in any subset of their domain), Milgrom and Weber (1982) show that the revenue equivalence theorem breaks down. In fact, in this situation, McAfee et al. (1989) show that the seller may extract the entire surplus from the bidders as

if there was no asymmetric information. Starting with Maskin and Riley (1989), several authors have also analyzed auctions of multiple units.

Finally, Clarke (1971) and Groves (1973) initiated another group of models in which the principal contracts with several agents simultaneously but does not attempt to maximize her own profits. This is the case of the provision of a public good through a mechanism provided by a benevolent regulator.

Relationships with Several Agents: Other Models and Organizational Design

Adverse selection models have attempted to analyze the optimal task assignment, the advantages of delegation, or the optimal structure of contractual relationships, when the principal contracts with several agents. Riordan and Sappington (1987) analyze a situation where two tasks have to be fulfilled and show that if the person in charge of each task has private information about the costs associated with the task, then the assignment of tasks within the organization is an important decision. For example, when the costs are positively correlated, then the principal will prefer to take charge of one of the phases herself, while she will prefer to delegate the task when the costs are negatively correlated.

In a very general framework, Myerson (1982) shows a powerful result: in adverse selection situations, centralization cannot be worse than decentralization, since it is always possible to replicate a decentralized contract with a centralized one. This result is really a generalization of the revelation principle. Baron and Besanko (1992) and Melumad et al. (1995) show that if the principal can offer complex contracts in a decentralized organization, then a decentralized structure can replicate a centralized organization. When there are problems of communication between principal and agents, the equivalence result does not hold: Melumad and Reichelstein (1987) show that *delegation* of authority can be preferable if communication between the principal and the agents is difficult. Still concerning the optimal design of the organization, Dana (1993) analyzes the optimal hierarchical structure in industries with several productive phases, when

firms have private information related to their costs. They show that structures that concentrate all tasks to a single agent are superior, since the incentives to dishonestly reveal the costs of each of the phases are weaker. Da-Rocha-Alvarez and De-Frutos (1999) argue that the absolute advantage of the centralized hierarchy is not maintained if the differences in costs between the different phases are sufficiently important.

Several Principals

Stole (1991) and Martimort (1996) point out the difficulty of extending the revelation principle to situations where an agent with private information is contracted by several principals who act separately. Given that not only one contract (or menu of contracts) is offered to the agent, but several contracts coming from different principals, it is not longer necessarily true that the best a principal can do is to offer a “truth-telling mechanism.”

Consider a situation with two principals that are hiring a single agent. If we accept that agent's messages are restricted to the set of possible types that the agent may have, we can obtain some conclusions. If the activities or efforts that the agent carries out for the two principals are substitutes (e.g., a firm producing for two different customers), then the usual result on the distortion of the decision holds: the most efficient type of agent supplies the efficient level of effort, while the effort demanded from the least efficient type is distorted. However, due to the lack of cooperation between principals, the distortion induced in the effort demanded from the less efficient type of agent is lower than the one maximizing the principals' aggregate profits. On the other hand, if the activities that the agent carries out for the principals are complementary (e.g., the firm produces a final good that requires two complementary intermediate goods in the production process), then the comparison of the results under cooperation and under no cooperation between the principals reveals that if a principal reduces the effort demanded from the agent, in the second case, this would imply that it is also profitable for the other principal to do the same. Therefore, the distortion in decisions is greater to that produced in the case in which principals cooperate.

Models of Moral Hazard and Adverse Selection

The analysis of principal-agent models where there are simultaneously elements of moral hazard and adverse selection is a complex extension of classic agency theory. Conclusions can be obtained only in particular scenarios. One common class of models considers situations where the principal cannot distinguish the part corresponding to effort from the part corresponding to the agent's efficiency characteristic because both variables determine the production level. Picard (1987) and Guesnerie et al. (1989) propose a model with risk-neutral participants and show that if the effort demanded from the different agents is not decreasing in characteristic (if a higher value of this parameter implies greater efficiency), then the optimal contract is a menu of distortionary deductibles designed to separate the agents. The menu of contracts includes one where the principal sells the firm to the agent (aiming at the most efficient type) and another contract where she sells only a part of the production at a lower prize (aiming at the least efficient type). However, there are also cases where fines are needed to induce the agents to honestly reveal their characteristic.

In fact, the main message of the previous literature is that the optimal solution for problems that mix adverse selection and moral hazard does not imply efficiency losses with respect to the pure adverse selection solution when the agent's effort is observable. However, in other frameworks (see Laffont and Tirole 1986), a true problem of asymmetric information appears only when both problems are mixed when, and efficiency losses are evident. Therefore, the same solution as when only the agent's characteristic is private information cannot be achieved.

hazard and adverse selection, where there is only one dimension in which information is asymmetric. A great deal of effort is devoted to try to ascertain whether it is moral hazard, or adverse selection, or both prevalent in the market. This is a difficult task because both adverse selection and moral hazard generate the same predictions in a cross section. For instance, a positive correlation between insurance coverage and probability of accident can be due to either the intrinsically riskier drivers selecting into contracts with better coverage (as the (Rothschild and Stiglitz 1976) model of adverse selection will predict) or to drivers with better coverage exerting less effort to drive carefully (as the canonical moral hazard model will predict). Chiappori et al. (2006) have shown that the positive correlation between coverage and risk holds more generally in the canonical models as long as the competitive assumption is maintained.

Future empirical approaches are likely to incorporate market power (as in Chiappori et al. 2006), multiple dimensions of asymmetric information (as in Finkelstein and McGarry 2006), as well as different measures of asymmetric information (as in Vera-Hernandez 2003). These advances will be partly possible thanks to richer surveys which collect subjective information regarding agents' attributes usually unobserved by principals or agent's subjective probability distributions. The wider availability of panel data will mean that it will become easier to disentangle moral hazard from adverse selection (as in Abbring et al. 2003). Much is to be learned by using field experiments that allow randomly varying contract characteristics offered to individuals and hence disentangling moral hazard from adverse selection (as in Karlan and Zinman 2009).

Future Directions

Empirical Studies of Principal-Agent Models

The growing interest on empirical issues related to asymmetric information started in the mid-1990s (see the survey by Chiappori and Salanie (2003)). A very large part of the literature is devoted to test the predictions of the canonical models of moral

Contracts and Social Preferences

Although principal-agent theory has proved fundamental in expanding our understanding of contract situations, real-life contracts frequently do not exactly match its predictions. Many contracts are linear and simpler, and incentives are often stronger and wage gaps more compressed than expected. One possible explanation is that the

theory has mainly focused on economic agents exclusively motivated by their own monetary incentives. However, this assumption leaves aside issues such as social ties, team spirit, or work morale, which the human resources literature highlights. A recent strand of economic literature, known as “behavioral contract theory,” has tried to incorporate social aspects into the economic modeling of contracts.

Such theory has been motivated by two types of empirical support. On the one hand, extensive interview studies with firm managers and employees (Bewley 1999) have shown not only that agents care about social comparisons such as internal pay equity or effort diversity but that their incentives to work hard are affected by them and that principals are aware of it and design their contracts accordingly. On the other hand, one of the most influential contributions of the experimental literature has been to show that, assuming that economic agents are not completely selfish (but exhibit some form of social preferences), it helps organizing many laboratory data. Experiments replicating labor markets (starting with Fehr (1993)) confirm Akerlof’s (1982) insight that contracts may be understood as a form of gift exchange in which principals may offer a “generous” wage and agents may respond with more than the minimum effort required.

Incorporating social and psychological aspects in a systematic manner into agents’ motivations has given rise to several forms of utility functions reflecting inequality aversion (Bolton and Ockenfels 2000; Desiraju and Sappington 2007; Fehr and Schmidt 1999), fairness (Rabin 1993), and reciprocity (Dufwenberg and Kirchsteiger 2004). More recently, such utility functions have been included into standard contract theory models and have helped in shortening the gap between theory predictions and real-life contracts. In particular, issues such as employees’ feelings of envy or guilt toward their bosses (Itoh 2004), utility comparisons among employees (Grund and Sliwka 2005; Rey-Biel 2008), or peer pressure motivating effort decisions (Huck and Rey-Biel 2006; Lazear 1995) have proved important in widening the scope of issues principal-agent theory can help to understand.

Principal-Agent Markets

The literature has been treating each principal-agent relation as an isolated entity. Thus, it normally takes a given relationship between a principal and an agent (or among several principals and/or several agents) and analyzes the optimal contract. In particular, the principal assumes all the bargaining power as she has the right to offer the contract she likes the most, and agent’s payoff is determined by his exogenously given reservation utility. However, in markets there is typically not a single partnership but there are several. It is then interesting to consider the simultaneous determination of both the identity of the pairs that meet (i.e., the *matching* between principals and agents) and the contracts these partnerships sign. The payoffs to each principal and agent will then depend on the other principal-agent relationships being formed in the market. This analysis requires a general equilibrium-like model.

Game theory provides a very useful tool to deal with the study of markets where heterogeneous players from one side can form partnerships with heterogeneous players from the other side: the two-sided matching models. Examples of classic situations studied in two-sided matching models (see Roth and Sotomayor 1990; Shapley and Shubik 1972) are the marriage market, the college admissions model, or the assignment market (where buyers and sellers transact). Several papers extend this game theory models to situations where each partnership involves contracts and show that the simultaneous consideration of matching and contracts has important implications. Dam and Pérez-Castrillo (2006) show that, in an economy where landowners contract with tenants, a government willing to improve the situation of the tenants can be interested in creating wealth asymmetries among them. Otherwise, the landowners would appropriate all the incremental money that the government is willing to provide to the agents. Serfes (2008) shows that higher-risk projects do not necessarily lead to lower incentives, which is the prediction in the standard principal-agent theory, and Alonso-Paulí and Pérez-Castrillo (2012) apply the theory to markets where contracts (between shareholders and managers) can include Codes of Best Practice. On the

empirical side, Akerberg and Botticini (2002) find strong evidence for endogenous matching between landlords and tenants and that risk sharing is an important determinant of contract choice.

Future research will extend the general equilibrium analysis of principal-agent contracts to other markets. In addition, the literature has only studied one-to-one matching models. This should be extended to situations where each principal can hire several agents or where each agent deals with several principals. The interplay between (external) market competition and (internal) collaboration between agents or principals can provide useful insights about the characteristics of optimal contracts in complex environments.

Bibliography

Primary Literature

- Abbring J, Chiappori PA, Heckman JJ, Pinquet J (2003) Adverse selection and moral hazard in insurance: can dynamic data help to distinguish? *J Eur Econ Assoc* 1:512–521
- Akerberg DA, Botticini M (2002) Endogenous matching and the empirical determinants of contract form. *J Pol Econ* 110:564–592
- Akerlof G (1970) The market for ‘Lemons’: qualitative uncertainty and the market mechanism. *Q J Econ* 89:488–500
- Akerlof G (1982) Labor contracts as a partial gift exchange. *Q J Econ* 97:543–569
- Alchian A, Demsetz H (1972) Production, information costs, and economic organization. *Am Econ Rev* 62:777–795
- Alonso-Pauli E, Pérez-Castrillo D (2012) Codes of best practice in competitive markets. *Econ Theory* 49(1):113–141
- Arrow K (1985) The economics of agency. In: Pratt J, Zeckhauser R (eds) *Principals and agents: the structure of business*. Harvard University Press, Boston
- Baron D, Besanko D (1984) Regulation and information in a continuing relationship. *Inf Econ Pol* 1:267–302
- Baron D, Besanko D (1987) Commitment and fairness in a dynamic regulatory relationship. *Rev Econ Stud* 54:413–436
- Baron D, Besanko D (1992) Information, control and organizational structure. *J Econ Manag Strategy* 1(2):237–275
- Baron D, Myerson R (1982) Regulating a monopoly with unknown costs. *Econometrica* 50:911–930
- Bernheim BD, Whinston MD (1986) Common agency. *Econometrica* 54:923–942
- Bewley T (1999) *Why rewards don’t fall during a recession*. Harvard University Press, Cambridge
- Bolton G, Ockenfels A (2000) ERC: a theory of equity, reciprocity and competition. *Am Econ Rev* 90:166–193
- Chiappori PA, Macho-Stadler I (1990) Contrats de Travail Répétés: Le Rôle de la Mémoire. *Ann Econ Stat* 17:4770
- Chiappori PA, Salanié B (2003) Testing contract theory: a survey of some recent work. In: Dewatripont H, Turnovsky (eds) *Advances in economics and econometrics*, vol 1. Cambridge University Press, Cambridge, pp 115–149
- Chiappori PA, Macho-Stadler I, Rey P, Salanié B (1994) Repeated moral hazard: the role of memory, commitment and access to credit markets. *Eur Econ Rev* 38:1527–1553
- Chiappori PA, Jullien B, Salanié B, Salanié F (2006) Asymmetric information in insurance: general testable implications. *Rand J Econ* 37:783–798
- Clarke E (1971) Multipart pricing of public goods. *Public Choice* 11:17–33
- Dam K, Pérez-Castrillo D (2006) The principal-agent matching market. *Frontiers Econ Theory Berkeley Electr* 2(1):1–34
- Dana JD (1993) The organization and scope of agents: regulating multiproduct industries. *J Econ Theory* 59:288–310
- Da-Rocha-Alvarez JM, De-Frutos MA (1999) A note on the optimal structure of production. *J Econ Theory* 89:234–246
- Demski JS, Sappington D (1986) Line-item reporting, factor acquisition and subcontracting. *J Account Res* 24:250–269
- Desiraju R, Sappington D (2007) Equity and adverse selection. *J Econ Manag Strategy* 16:285–318
- Dufwenberg M, Kirchsteiger G (2004) A theory of sequential reciprocity. *Games Econ Behav* 47:268–298
- Fehr E, Schmidt K (1999) A theory of fairness, competition and cooperation. *Q J Econ* 114:817–868
- Fehr E, Kirchsteiger G, Riedl A (1993) Does fairness prevent market clearing? *Q J Econ* 108:437–460
- Finkelstein A, McGarry K (2006) Multiple dimensions of private information: evidence from the long-term care insurance market. *Am Econ Rev* 96:938–958
- Freixas X, Guesnerie R, Tirole J (1985) Planning under information and the ratchet effect. *Rev Econ Stud* 52:173–192
- Fudenberg D, Holmström B, Milgrom B (1990) Short-term contracts and long-term agency relationships. *J Econ Theory* 51:1–31
- Gibbard A (1973) Manipulation for voting schemes. *Econometrica* 41:587–601
- Green JR, Laffont JJ (1977) Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica* 45:427–438
- Grossman SJ, Hart OD (1983) An analysis of the principal-agent problem. *Econometrica* 51:7–45
- Groves T (1973) Incentives in teams. *Econometrica* 41:617–631
- Grund C, Sliwka D (2005) Envy and compassion in tournaments. *J Econ Manag Strateg* 14:187–207

- Guesnerie R, Laffont JJ (1984) A complete solution to a class of principal-agent problems with application to the control of a self-managed firm. *J Public Econ* 25:329–369
- Guesnerie R, Picard P, Rey P (1989) Adverse selection and moral hazard with risk neutral agents. *Eur Econ Rev* 33:807–823
- Harris M, Raviv A (1978) Some results on incentive contracts with applications to education and employment, health insurance and law enforcement. *Am Econ Rev* 68:20–30
- Harris M, Raviv A (1979) Optimal incentive contracts with imperfect information. *J Econ Theory* 2:231–259
- Holmström B (1979) Moral hazard and observability. *Bell J Econ* 10:74–91
- Holmström B (1982) Moral hazard in teams. *Bell J Econ* 13:324–340
- Holmström B, Milgrom P (1991) Multitask principal-agent analysis: incentive contracts, assets ownership, and job design. *J Law Econ Organ* 7(Suppl):24–52
- Huck S, Rey-Biel P (2006) Endogenous leadership in teams. *J Inst Theor Econ* 162:1–9
- Itoh H (1990) Incentives to help in multi-agent situations. *Econometrica* 59:611–636
- Itoh H (2004) Moral hazard and other-regarding preferences. *Jpn Econ Rev* 55:18–45
- Jensen M, Meckling W (1976) The theory of the firm, managerial behavior, agency costs and ownership structure. *J Finan Econ* 3:305–360
- Karlan D, Zinman J (2009) Observing unobservables: identifying information asymmetries with a consumer credit field experiment. *Econometrica* 77:1993–2008
- Klemperer P (2004) Auctions: theory and practice. Princeton University Press, Princeton
- Laffont JJ (1990) Analysis of hidden gaming in a three levels hierarchy. *J Law Econ Organ* 6:301–324
- Laffont JJ, Tirole J (1986) Using cost observation to regulate firms. *J Polit Econ* 94:614–641
- Laffont JJ, Tirole J (1987) Comparative statics of the optimal dynamic incentive contract. *Eur Econ Rev* 31:901–926
- Laffont JJ, Tirole J (1988) The dynamics of incentive contracts. *Econometrica* 56:1153–1176
- Laffont JJ, Tirole J (1990) Adverse selection and renegotiation in procurement. *Rev Econ Stud* 57:597–626
- Lambert R (1983) Long term contracts and moral hazard. *Bell J Econ* 14:441–452
- Lazear E (1995) Personnel economics. MIT Press, Cambridge
- Macho-Stadler I, Pérez-Castrillo D (1993) Moral hazard with several agents: the gains from cooperation. *Int J Ind Organ* 11:73–100
- Macho-Stadler I, Pérez-Castrillo D (1998) Centralized and decentralized contracts in a moral hazard environment. *J Ind Econ* 46:489–510
- Malcomson JM, Spinnnewyn F (1988) The multiperiod principal-agent problem. *Rev Econ Stud* 55:391–408
- Martimort D (1996) Exclusive dealing, common agency and multiprincipals incentive theory. *Rand J Econ* 27:1–31
- Maskin E, Riley J (1989) Optimal multi-unit auctions. In: Hahn F (ed) *The economics of missing markets, information, and games*. Oxford University Press, Oxford, pp 312–335
- McAfee P, McMillan J, Reny P (1989) Extracting the surplus in the common value auction. *Econometrica* 5:1451–1459
- Melumad N, Reichelstein S (1987) Centralization vs delegation and the value of communication. *J Account Res* 25:1–18
- Melumad N, Mookherjee D, Reichelstein S (1995) Hierarchical decentralization of incentive contracts. *Rand J Econ* 26:654–672
- Milgrom P, Weber RJ (1982) A theory of auctions and competitive bidding. *Econometrica* 50:1089–1122
- Mirrlees J (1971) An exploration in the theory of optimum income taxation. *Rev Econ Stud* 38:175–208
- Mirrlees J (1975) The theory of moral hazard and unobservable behavior, part I. WP Nuffield College, Oxford
- Mookherjee D (1984) Optimal incentive schemes with many agents. *Rev Econ Stud* 51:433–446
- Mussa M, Rosen S (1978) Monopoly and product quality. *J Econ Theory* 18:301–317
- Myerson R (1979) Incentive compatibility and the bargaining problem. *Econometrica* 47:61–73
- Myerson R (1981) Optimal auction design. *Math Oper Res* 6:58–73
- Myerson R (1982) Optimal coordination mechanisms in generalized principal-agent models. *J Math Econ* 10:67–81
- Pauly MV (1968) The economics of moral hazard. *Am Econ Rev* 58:531–537
- Pauly MV (1974) Overinsurance and public provision of insurance: the roles of moral hazard and adverse selection. *Q J Econ* 88:44–62
- Picard P (1987) On the design of incentive schemes under moral hazard and adverse selection. *J Public Econ* 33:305–331
- Rabin M (1993) Incorporating fairness into game theory and economics. *Am Econ Rev* 83:1281–1302
- Radner R (1981) Monitoring cooperative agreements in a repeated principal-agent relationship. *Econometrica* 49:1127–1148
- Rey P, Salanié B (1990) Long term, short term and renegotiation. *Econometrica* 58:597–619
- Rey-Biel P (2008) Inequity aversion and team incentives. *ELSE WP Scand J Econ* 110:297–320
- Riordan MH, Sappington DE (1987) Information, incentives, and the organizational mode. *Q J Econ* 102:243–263
- Rogerson W (1985) Repeated moral hazard. *Econometrica* 53:69–76
- Ross SA (1973) The economic theory of agency: the principal's problem. *Am Econ Rev* 63:134–139

- Roth AE, Sotomayor M (1990) Two-sided matching: a study in game-theoretic modeling and analysis. Cambridge University Press, New York
- Rothschild M, Stiglitz J (1976) Equilibrium in competitive insurance markets: an essay in the economics of imperfect information. *Q J Econ* 90:629–650
- Rubinstein A, Yaari ME (1983) Repeated insurance contracts and moral hazard. *J Econ Theory* 30:74–97
- Serfes K (2008) Endogenous matching in a market with heterogeneous principals and agents. *Int J Game Theory* 36:587–619
- Shapley LS, Shubik M (1972) The assignment game I: the core. *Int J Game Theory* 1:111–130
- Spence M (1974) Market signaling. Harvard University Press, Cambridge
- Spence M, Zeckhauser R (1971) Insurance information, and individual action. *Am Econ Rev* 61:380–387
- Stole L (1991) Mechanism design under common agency. WP MIT, Cambridge
- Tirole J (1986) Hierarchies and bureaucracies: on the role of collusion in organizations. *J Law Econ Organ* 2:181–214
- Vera-Hernandez M (2003) Structural estimation of a principal-agent model: moral hazard in medical insurance. *Rand J Econ* 34:670–693
- Vickrey W (1961) Counterspeculation, auctions and competitive sealed tenders. *J Finan* 16:8–37
- Vickrey W (1962) Auction and bidding games. In: Recent advances in game theory. The Princeton University Conference Proceedings, Princeton, New Jersey, pp 15–27

Books and Reviews

- Hart O, Holmström B (1987) The theory of contracts. In: Bewley T (ed) *Advances in economic theory*, fifth world congress. Cambridge University Press, Cambridge
- Hirshleifer J, Riley JG (1992) *The analytics of uncertainty and information*. Cambridge University Press, Cambridge
- Laffont JJ, Martimort D (2002) *The theory of incentives: the principal-agent model*. Princeton University Press, Princeton
- Laffont JJ, Tirole J (1993) *A theory of incentives in procurement and regulation*. MIT Press, Cambridge
- Macho-Stadler I, Pérez-Castrillo D (1997) *An introduction to the economics of information: incentives and contracts*. Oxford University Press, Oxford
- Milgrom P, Roberts J (1992) *Economics, organization and management*. Prentice-Hall, Englewood Cliffs

Differential Games

Marc Quincampoix

Laboratoire de Mathématiques de Bretagne
Atlantique (LMBA), Université de Brest, Brest,
France

Article Outline

Glossary

Definition of the Subject and Its Importance

Introduction

Qualitative and Quantitative Differential Games

Existence of a Value for Zero Sum Differential
Games

Nonantagonist Differential Games

Stochastic Differential Games

Differential Games with Incomplete Information

Miscellaneous

Bibliography

Glossary

Dynamics This is the law which governs the evolution of the system: for differential games it is a differential equation.

Strategies This is the way a player chooses his control as a function of the state of the system and of the action of his opponents.

Information This is the set of parameters known by the player in order to build his strategy.

Definition of the Subject and Its Importance

Differential games is a mathematical theory which concerns with problems of conflicts modeled as game problems in which the state of the players depends on the time in a continuous way. The positions of the players are solution to differential equations. Differential games can be described from two different points of views, depending

mainly of the field of applications. Firstly, they can be considered as games where the time is continuous. This aspect is often considered for applications in economics or management sciences. Secondly, they also can be viewed as control problems with several controller having different objectives. In this way, differential games are a part of control theory with conflicts between the players. The second aspect concerns often classical applications of control theory: engineers sciences.

The importance of the subject was emphasized by J. Von Neuman in 1946 in his pioneer book “Theory of games and Economic Behaviour” (Von Neumann and Morgenstern 1946) *We repeat most emphatically that our theory is thoroughly static. A dynamic theory would unquestionably be more complete and therefore preferable.* But at this date, the main efforts were devoted to treat static aspects of game theory. The true birth of the domain was pursuit differential games (motivated by military applications in the “Cold War”) developed in the 1950s concurrently developed by R. Isaacs in USA and by L. Pontryagin in Soviet Union (Pontryagin 1968). Now differential games have a wide range of applications from Economics to engineers sciences but also more recently to biology, behavioral ecology and population dynamics. The present article is focused on two-player zero sum and antagonist differential games.

Introduction

The best-known example of differential game is the pursuit game “Lion and Man” where in a infinite plane a Lion want to catch a Man. This example is elementary enough for describing clearly the problematic and we will use this example as an illustration throughout the article. The position of the Man at time t is $y(t)$ while the Lion position is $z(t)$. At any time, the Man can choose his velocity $y'(t)$ in any direction but with a maximum intensity, while the Lion can choose his velocity $z'(t)$ with an intensity less or equal than

L (here $L > M$). The objective of the Lion is to catch the Man as soon as possible; the aim of the Man is to escape to the Lion as long as possible. In this very intuitive game, we wish to introduce the important elements needed for defining a differential game: *the dynamics, the actions of the players, the objectives, the rules of the games (the strategies)*. The *dynamics* is then

$$\begin{aligned} y'(t) &= u(t), \quad u(t) \in U, \quad z'(t) = v(t) \in V; \quad t \\ &\geq 0 \end{aligned} \quad (1)$$

where $u(t)$ is the *action* chosen by the first player (the Man), $u(t)$ belongs to a set $U = \{u \mid \|u\| \leq M\}$ similarly $v(t)$ belongs to the set $V = \{v \mid \|v\| \leq L\}$. The *objectives* of the two players are of antagonist nature: The Lion wants to minimize the time before capture while the Man wants to maximize it. The rule of the game is the way the two players choose their actions: here we can assume that they choose their instantaneous velocity as a function of the current positions

$$u(t) = \bar{u}(x(t), y(t)), \quad v(t) = \bar{v}(x(t), y(t)) \quad (2)$$

Clearly the time of capture T in a function of the strategies $\bar{u}$ and $\bar{v}$; the Lion wants to minimize this function, the Man wants to maximize it. Solving the game means to find the *optimal value*, namely, the time of capture which is minimum over all strategies $\bar{v}$ and maximum over all strategies $\bar{u}$. In the elementary example of “Lion of Man,” one can easily check that

$$\begin{aligned} \bar{v}(y, z) &= L \frac{y - z}{\|y - z\|}, \quad \bar{u}(y, z) = -M \frac{y - z}{\|y - z\|}, \\ T &= \frac{\|y(0) - z(0)\|}{L - M}, \end{aligned}$$

which is a rather intuitive solution: the Lion runs in direction to the Man at his maximal possible speed L , while the Man runs in the opposite direction with his maximal speed M . We leave this illustrative example which has allowed to underline the concepts of differential games, we will present now.

The present article is mainly restricted to the two players case. A state variable x is governed by a differential equation

$$x'(t) = f(x(t), u(t), v(t)). \quad (3)$$

The players act on the state variable x by choosing two controls u and v which are functions of the time variable t taking their values in spaces U and V . (Note that for the Lion of the Man $x = (y, z)$ and each player controls a part of the state: the games is said *separated*.) For making the notations clearer, the first player which chooses u is called Ursula and the second player which chooses v is called Victor.

A classification of differential games according to the level of conflict between the players is very relevant. The case of opposite objectives of the two players is concerned by *antagonist games* or *zero-sum games*, while the case of non opposite objective is concerned by *cooperative or non-antagonist* differential games.

The objectives of the game can be quantitative – players want to minimize or maximize a *payoff* (a function depending on the initial state and their action, for instance the time of capture in the Lion an Man case) – or qualitative for instance one player wants the state to stay in a subset of the state space while the other player wants the state to reach a given target. In Isaacs’ terminology, these games are *games of degree* or *games of kind*. The definition of the rules of the games (or strategies) is one of the most difficult topics in differential games; this problem does not exist for static games or for elementary pursuit games like Lion and Man. To illustrate this fact, suppose that we have solve the game with a strategy of *feedback form*, i.e., when controls depend on the current state (as in Eq. (2)), then we have to solve the differential equation

$$x'(t) = f(x(t), \bar{u}(x(t)), \bar{v}(x(t)))$$

which can has neither uniqueness nor existence of solutions, even if Eq. (3) enjoys good properties of existence and uniqueness of solutions. This point is crucial and is of pure mathematical nature. A way to overcome this difficulty could be to impose regularity properties of the feedback strategies. In the – young – history of differential games, this was the “time of paradoxes” in the 1960s–1970s, where specific regularity of

feedback was adequate for a class of problem but lead to “paradoxes” (for quantitative game nonexistence of the value, for qualitative games existence of a initial position starting from which players are not able neither to win nor to lose) in slightly modified problems. This was solved in the 1970s by enlarging the class of feedback strategies allowing the choice of the control to depend not only on the current state space but also on passed values of the control of the opponents (this is the so-called nonanticipative strategies introduced by Varaiya Roxin Elliot Kalton (Varaiya 1967; Varaiya and Lin 1967; Roxin 1969; Elliot and Kalton 1972), another class of strategies was introduced par Krasovski (Breakwell 1977).

For differential games, another important feature is to express the fact that the two players acts simultaneously on the system. One way to translate this fact in a mathematical way is to prove that the order of the actions of the player does not modify the final result. Section “[Qualitative and Quantitative Differential Games](#)” will concerns with this question. For instance, in the case of quantitative games like Lion and Man, the question is that interchanging the order of operations “minimum” and “maximum” we must obtain the same result. This is the problem of existence of the value which is of first importance for differential games. We discuss this question in section “[Existence of a Value for Zero Sum Differential Games](#).” The fifth section is devoted to cooperative quantitative games. Some other aspect or applications of differential games are evoked in section “[Stochastic Differential Games](#).” In the last section, we give short descriptions of very actual domains of researches as problems of information or impulsive differential games.

Qualitative and Quantitative Differential Games

Throughout this article, we will make suppositions such that as soon as the initial position is fixed and the controls u and v are given, we have the existence of a unique associated trajectory defined for every time (this could be easily obtained with assumptions on the function f).

We also denote X the state space to which the state variable x belongs.

A nonanticipative strategy for the first player associates to any action of the second player (to any control v), a control u of the first player such that *at any time the v depends only of past values of v* . It could be shown that every regular feedback strategy is such a nonanticipative strategy. So a nonanticipative strategy of the first player is a map α from $\mathcal{V}$ (the set of measurable control $v[0, +\infty) \mapsto V$) on $\mathcal{U}$ (the set of measurable controls u), which has furthermore the non-anticipative property. Similarly, one can define a strategy of the second player as a nonanticipative function β from $\mathcal{U}$ to $\mathcal{V}$.

We will now present a rather general description of a target game viewed as a game of kind (qualitative game) or as a game of degree. For this, we consider a given set in the state space – called the target – that Victor wants the state of the system to reach while Ursula wants to avoid the target. For the Lion and Man pursuit game, the target can be considered as the set of (y, z) such that $y = z$.

Qualitative Target Games

The problems consists in finding the victory domains, i.e., the set of initial positions such that one player can win. This leads to the precise following definition.

Victor Victory domain W_V is the set of initial position x_0 (not in C) such that he can find a strategy β such that there exists a time T such that whatever is the control u chosen by Ursula, the target is reached before the time T (cf. Aubin 1991).

In fact, to make this definition mathematically correct, we have to add a small number ε and to say that the trajectory reaches a ε neighborhood of C :

$$W_V := \{x_0 \notin C, \exists \beta, \exists T, \varepsilon > 0, \text{ such that } \forall u \in \mathcal{U}, \\ \exists \tau > 0, \text{ dist}(x[x_0, u, \beta(u)](\tau), C) \leq \varepsilon\}.$$

In the above formula, dist means the distance and $t \mapsto x[x_0, u, \beta(u)](t)$ denotes the trajectory associated with the control u and the strategy β .

In a parallel way, the victory domain W_U of Ursula is the set of initial positions x_0 for which

Ursula can find a strategy α such that whatever is the control played by Victor, the associate trajectory never reaches the target C .

$$W_U := \{x_0 \notin C, \exists \alpha, \text{ such that } \forall v \in \mathcal{V}, \\ x[x_0, \alpha(v), v](t) \notin C, \forall t \geq 0\}.$$

Before going further, it could be surprising to the reader that the game is not “symmetric” due to the presence of $\varepsilon > 0$. This is a mathematical problem that exceed the scope of the present paper, for better understanding the reader can imagine that $\varepsilon = 0$ and see Cardaliaguet (1996) for more deep analysis.

The main problem of such a qualitative game is the following alternative problem. Roughly speaking, if one player does not win, the other player must win. Of course this appears to be the minimal requirement in order the modelization of the problem is correctly formulated. Nevertheless, the alternative is not an easy problem; for instance, it is not obvious that the intuitive notions of feedback strategies are not suitable for giving a positive answer of this question. The alternative can be expressed by the fact that the sets C , W_U and W_V form a partition of the whole space X :

$$X = C \cup W_U \cup W_V, \text{ and } W_U \cap W_V = \emptyset.$$

The above alternative is valid under several technical conditions (for instance C is an open set); we will not describe now and under the following crucial condition – called *Isaacs condition*

$$\min_{v \in V} \max_{u \in U} \langle f(x, u, v), p \rangle = \max_{u \in U} \min_{v \in V} \langle f(x, u, v), p \rangle \quad (4)$$

for any direction p . This could be understood as the existence of a saddle point for the *static game* with payoff $\langle f(x, u, v), p \rangle$ ($\langle \cdot, \cdot \rangle$ is the scalar product). The expression of the equation (4) above is called the *Hamiltonian* of the game; it is a function $(x, p) \mapsto H(x, p)$.

The alternative theorem was firstly obtained by Krassovskii and Subbotin for a slightly different notion of strategy, and by Cardaliaguet (1996) for nonanticipative strategy. It is worth pointing out that even for the elementary Lion and Man game

played in a circular “arena” (instead of the whole plane), the problem is not obvious (Flynn 1973). For more complex forms of “arenas” and for general differential games with restricted space domain, the alternative problem was only solved in 2001 (Cardaliaguet et al. 2001).

Once the alternative property is known, the second interesting step is to describe the victory domains. In several examples, Isaacs has discovered that the boundaries of the domains satisfy a geometric equation called *the Isaacs equations*

$$H(x, v_x) = 0 \quad \text{for any } x \text{ on the boundary} \\ \text{of the victory domain} \quad (5)$$

where v_x denotes the normal of the victory domain. Hypersurfaces satisfying such equation are *semipermeable barriers*; one player can prevent the other player to make the state crossing the barrier. From this property, it is possible to get many information on the winning strategies. But unfortunately, the victory domains are seldom regular enough to have a normal. Very often they have “corners” even for very elementary games. So there are two different ways to treat this question. A first approach initiated by Isaacs himself, and developed by Breakwell, Bernhard, Melykian consists in a precise and fine geometrical analysis of semipermeable barriers (Breakwell 1977; Bernhard 1988; Buckdahn et al. 2009a). When this approach is possible, this gives very precise information on the behavior of the players. Unfortunately, this is hardly possible in high dimension games and/or when the victory domain is not smooth enough. A second approach consists in proving that the boundary of the victory domain satisfies (5) in suitable generalized sense (Quincampoix 1992; Cardaliaguet 1997) and to use this information to approach numerically the victory domain (cf. Cardaliaguet et al. (1999) for a detailed exposition of this method).

Quantitative Target Games

Here the goal of the player is of quantitative nature: Victor wants to minimize the time to reach the target C while Ursula wants to minimize it. For describing the game, we associate to any

trajectory of the dynamics $t \in [0, +\infty) \mapsto x(t)$ the first time $x(\cdot)$ reaches C :

$$\vartheta_C(x(\cdot)) = \inf\{t \geq 0 \mid x(t) \in C\},$$

with, by convention, $\vartheta_C(x(\cdot)) = +\infty$ if $x(\cdot)$ does not reach C .

The modelization problem of such a quantitative game is to express the fact that both players act *simultaneously* on the system. So roughly speaking, we must check that if Ursula chooses her strategy before the second player, the result is the same that in the case when Victor chooses his strategy first. This means that the following *value functions* of the game do coincide.

$$\begin{aligned}\vartheta^b(x_0) &= \inf_{\beta} \sup_{u \in \mathcal{U}} \vartheta_C(x[x_0, u, \beta(u)]) \quad (\text{lower value}), \\ \vartheta^\#(x_0) &= \sup_{\alpha} \inf_{v \in \mathcal{V}} \vartheta_C(x[x_0, \alpha(v), v]) \quad (\text{upper value}).\end{aligned}$$

When $\vartheta^b = \vartheta^\#$, one says that the game has a value. Because the question of existence of value in differential game is an essential question in game theory, the section “[Existence of a Value for Zero Sum Differential Games](#)” is devoted to the exposition of this feature.

It is worth pointing out, that oppositely to many static games, this quantitative differential game is not in a *normal form*, namely, a player does not play a strategy against a strategy of his opponents. This is another main difference with static games. In fact, it is sometimes possible to write the game in normal form allowing the non-anticipative strategies to have a small delay (Cardaliaguet and Quincampoix 2008). The normal form will be used in section “[Nonantagonist Differential Games](#)” for nonzero sum games.

The *Dynamic programming principle* says, roughly speaking, that the game maintains the same structure if the time change. Indeed if starting from time 0 the players plays the games until time t , then at time t both players are facing to a differential game of the same nature than the initial games. This can be expressed by the dynamic programming equations

$$\begin{cases} \vartheta^b(x_0) = \inf_{\beta} \sup_{u \in \mathcal{U}} \{t + \vartheta^b(x[x_0, u, \beta(u)](t))\} \\ \vartheta^\#(x_0) = \sup_{\alpha} \inf_{v \in \mathcal{V}} \{t + \vartheta^\#(x[x_0, \alpha(v), v](t))\} \end{cases} \quad (6)$$

which is available as soon as the trajectories do not reach the target.

Existence of a Value for Zero Sum Differential Games

Using the dynamic programming principle, it is possible to deduce an infinitesimal characterization of the value functions (formally we subtract the two sides of one line of (6), divide by t and we let t tend to 0^+). Under the Isaacs condition (4), when the value functions are differentiable, it is not difficult to show that the values $\vartheta^\#$ and ϑ^b are solutions to a partial differential equation: the following *Hamilton Jacobi Isaacs Equation*

$$H(x, D\vartheta(x)) = 1 \quad \text{in } \mathbb{R}^N \setminus C. \quad (7)$$

This fact was noticed in the very beginning of the history of differential games. Furthermore, if the Hamilton Jacobi Equation has a continuously differentiable solution, then this solution is the value function and the game can be solved using feedback strategies (for instance, the Lion and Man game). This is the famous *Isaacs verification Theorem*. This is a completely satisfactory solvation of the game when the value is smooth. Unfortunately, very early it was noticed that the values are not differentiable and so the previous approach is not possible.

Before going further, it is worth pointing out that this is not only a mathematical question of regularity of the value functions. Indeed the smoothness of the value function is closely related of the number of optimal strategies in the game. The most of differential games, even very elementary ones, have not smooth values. The reader can convince himself by considering the Lion of the Man game in a circular arena with furthermore a round pillar in the center of the arena (that of course the players cannot cross).

One of the more substantial progresses in the theory of differential game in the 1980s was due to the use of viscosity solution theory (Bardi and Capuzzo Dolcetta 1997). In fact, even when the values $\vartheta^\#$ and ϑ^b are not smooth, they are both solutions in *viscosity sense* to the Hamilton Jacobi Isaacs Equation (7). Moreover with very weak assumptions (Lipschitzian continuity of the data), the partial differential equation (7) has a unique solution. Consequently both values coincide. This existence of the value result was due to Evans and Souganidis (1984) and can be viewed as a generalization of Isaacs verification theorem. We refer the reader to Bardi and Capuzzo Dolcetta (1997) for viscosity solutions. In several cases, for instance, in state constraint games case, the value is neither smooth nor Lipschitz (even not continuous), so it is necessary to extend the Evans Souganidis scheme. The main idea is to reduce the quantitative game to a qualitative game in a higher dimensional space (Cardaliaguet et al. 2001; Bettiol et al. 2006; Subbotin 1995; Berkovitz 1994); this is the object of the *Viability theory* approach to differential games. We refer the reader to Buckdahn et al. (2004), pp. 3–37 for a survey of this method. Another advantage of this *reduction to qualitative* method lies in the fact that numerical analysis tools of the qualitative approach (Cardaliaguet et al. 1999) can be used to build algorithms computing the value.

At this point, it is important to note that the scheme for obtaining the existence of the value is valid for a wide class of payoff. This fact can be explained by considering a game on a finite time interval $[0, T]$ such that at any initial condition $x(t_0) = x_0$ and any pair of control u and v is associated the cost

$$C(t_0, x_0, uv) := \int_{t_0}^T L(x(s), u(s), v(s)) ds + g(x(T)) \quad (8)$$

where L and g are given and $x(\cdot)$ is the solution to (3) on $[t_0, T]$ with initial condition $x(t_0) = x_0$. The payoff (the cost) is then the sum of an integral cost and of a terminal cost. Victor wants to minimize this cost while Ursula wants to maximize it. This leads to the following values

$$V^b(t_0, x_0) := \inf_{\beta} \sup_{u \in \mathcal{U}} C(t_0, x_0, u, \beta(u)),$$

$$V^\#(t_0, x_0) := \sup_{\alpha} \inf_{v \in \mathcal{V}} C(t_0, x_0, \alpha(v), v)$$

which are both solution to the following Hamilton Jacobi Isaacs equation

$$\begin{cases} \frac{\partial V}{\partial t}(t, x) + H\left(x, \frac{\partial V}{\partial x}(t, x)\right) = 0 & \text{for all } (t, x) \in [0, T] \times \mathbb{R}^N \\ V(T, x) = g(x) & \text{for all } x \in \mathbb{R}^N \end{cases} \quad (9)$$

Here the Isaacs condition takes the following form

$$\begin{aligned} \min_{v \in V} \max_{u \in U} \{L(x, u, v) + \langle f(x, u, v), \xi \rangle\} \\ = \max_{u \in U} \min_{v \in V} \{L(x, u, v) + \langle f(x, u, v), \xi \rangle\} \end{aligned}$$

for all $\xi \in \mathbb{R}^N$ and the Hamiltonian $H(x, \xi)$ is equal to the above expression (we refer to Plaskacz and Quincampoix (2000) and Berkovitz (1994) in the discontinuous or constrained case). Since the Hamilton Jacobi equations has at most one viscosity solution (Evans and Souganidis 1984), then $V^b = V^\#$ and so the game has a value.

This way of reasoning strongly depends on the Isaacs' condition. When the Isaacs' condition is not valid, the value does not exist in general with pure nonanticipative strategies, but it is possible to prove the existence of the value with suitable random strategies (cf. Buckdahn et al. (2013) and also Buckdahn et al. (2016) for a simplified version of mixed strategies and Buckdahn et al. (2014) for stochastic differential games).

Worst Case Design

An powerful application of antagonist games concerns zero-sum quantitative game where a controller wants to drive a system against another action on the system (uncertainty, nature) considered as a second player. The controller wants to find a *robust* strategy available against the worst case of the nature. Clearly the question of existence of a value is not relevant in this case. Only one value is interesting in this case.

The linear quadratic case was extensively studied in the book (Basar and Bernhard 1995). In the fully nonlinear case, the Hamilton Jacobi Isaacs Equation becomes infinite dimensional. An alternative approach can be done by setting Hamilton Jacobi Equations on the space of closed sets (Quincampoix and Veliov 2005).

Impulsive Games

Impulsive differential games where the player can at any time either follow a continuous dynamics of the form (3) or make a jump. These games combine the effects and difficulties of static and differential games. Only very preliminaries result are available (Buckdahn et al. 2004, pp. 223–249) for pursuit games when both players cannot “jump” simultaneously. Another interesting classes of pursuit impulsive game is the case when each player can choose at every time between two dynamics. This domain is still widely open.

Nonantagonist Differential Games

This section is devoted to a two player game on $[0, T]$ with different objectives: Ursula wants to maximize a cost

$$C_1(t_0, x_0, u, v) := \int_{t_0}^T L_1(x(s), u(s), v(s)) ds + g_1(x(T)),$$

while Victor wants to maximize a payoff

$$C_2(t_0, x_0, u, v) := \int_{t_0}^T L_2(x(s), u(s), v(s)) ds + g_2(x(T)).$$

Observe that if $C_2 + C_1 = 0$ (i.e., $L_1 + L_2 = 0$ and $g_1 + g_2 = 0$), the game reduces to the game studied in section “Existence of a Value for Zero Sum Differential Games.” Here we consider the general case where $C_1 + C_2$ is non necessarily equal to 0. In a nonantagonist game, it is important to play a strategy against a strategy, so we introduced the concept of nonanticipative strategies with delay on the time interval $[0, T]$. A nonanticipative strategy with delay for the first player associates to any action of the second player (to any control v) a control u of the first player such that there exists a delay $r > 0$

such that at any time t if two controls v_1 and v_2 coincide on $[0, t]$, then the associated controls $u_1 = \alpha(v_1)$ and $u_2 = \alpha(v_2)$ do coincide on $[0, t + r]$. The strategy nonanticipative with delay β for the second player is defined in a similar way. Clearly any strategy nonanticipative with delay is nonanticipative. Moreover, it is possible to prove (Cardaliaguet and Quincampoix 2008) that for a fixed initial condition (t_0, x_0) and for any pair (α, β) of nonanticipative strategies with delays, there exists a *unique pair* of controls (u, v) satisfying

$$\alpha(v) = u, \quad \beta(u) = v.$$

Hence it is possible to associate a trajectory to (α, β) which is the trajectory associated to (u, v) and we denote $C_1(t_0, x_0, \alpha, \beta)$ and $C_2(t_0, x_0, \alpha, \beta)$ the associated costs. This enables to write the game in a normal form.

The antagonist differential game problem consists in finding *Nash Equilibria* defined as follows. Fix (t_0, x_0) an initial condition, a pair of real numbers (e_1, e_2) is a Nash equilibria payoff if and only if there exists a pair of nonanticipative strategies with delays $(\bar{\alpha}, \bar{\beta})$ such that

$$e_1 = C_1(t_0, x_0, \bar{\alpha}, \bar{\beta}), \quad e_2 = C_2(t_0, x_0, \bar{\alpha}, \bar{\beta})$$

and such that for any other pair of nonanticipative strategy with delay (α, β) , the following inequalities hold true

$$\begin{aligned} C_1(t_0, x_0, \bar{\alpha}, \bar{\beta}) &\geq C_1(t_0, x_0, \alpha, \bar{\beta}), \\ C_2(t_0, x_0, \bar{\alpha}, \bar{\beta}) &\geq C_2(t_0, x_0, \bar{\alpha}, \beta). \end{aligned}$$

In a completely rigorous mathematical point of view, the above definition must be understood with a small $\varepsilon > 0$ error (the correct statement is: For all $\varepsilon > 0$, there exists $(\bar{\alpha}, \bar{\beta})$ such that

$$|e_i - C_i(t_0, x_0, \bar{\alpha}, \bar{\beta})| \leq \varepsilon, \quad i = 1, 2$$

and for any other pair of strategies (α, β) we have

$$\begin{aligned} C_1(t_0, x_0, \bar{\alpha}, \bar{\beta}) &\geq C_1(t_0, x_0, \alpha, \bar{\beta}) - \varepsilon, \\ C_2(t_0, x_0, \bar{\alpha}, \bar{\beta}) &\geq C_2(t_0, x_0, \bar{\alpha}, \beta) - \varepsilon \end{aligned}$$

cf. Cardaliaguet and Plasckacz (2003). Nash equilibria were studied in detail in Klejmenov (1993)

for another concept of strategy. Now the remaining part of this section concerns first the characterization of the Nash equilibria through zero-sum games and second the existence of Nash equilibria. For doing this, we introduce the following value function associated with two auxiliary zero-sum games where one player is the minimizer while its opponent is the maximizer.

$$V_1(t_0, x_0) = \inf_{\beta} \sup_{u \in \mathcal{U}(t_0)} C_1(t_0, x_0, u, \beta(u)),$$

$$V_2(t_0, x_0) = \inf_{\alpha} \sup_{v \in \mathcal{V}(t_0)} C_2(t_0, x_0, \alpha(v), v).$$

Recall that under Isaacs condition, the value of above auxiliary games does exist. Then now the characterization of Nash payoff can be expressed: the pair (e_1, e_2) is a Nash equilibrium payoff if and only if there exist two controls (u, v) such that

$$e_1 = C(t_0, x_0, u, v), \quad e_2 = C(t_0, x_0, u, v)$$

and for all time $t \in [t_0, T]$

$$e_1 \geq V_1(t, x[t_0, x_0, u, v](t)),$$

$$e_2 \geq V_2(t, x[t_0, x_0, u, v](t))$$

Observe that this is not an intuitive result (Once again a correct statement is, a $\epsilon > 0$ is needed to have a rigorous statement: For any ϵ there exists a pair (u, v) satisfying above inequality up to an error ϵ). From this result, it is possible – but not easy – to obtain the existence of a Nash equilibrium pair under Isaacs condition.

Observe that at a first glance, it is very surprising to obtain the existence of the Nash equilibria assuming Isaacs condition which means the existence of saddle point of static games. One could assume instead the existence of Nash equilibria of static games. Unfortunately, except for very specific games like some linear quadratic differential games, this method fails. Another way to understand this lies in the fact that Nash equilibria are very instable (cf. Buckdahn et al. 2004, pp. 57–67). One can also try to prove the existence of equilibria in feedback strategies; it is not always possible and even if one such equilibrium exists, it is very

unstable (Bressan 2011). We refer the reader to the bibliography for noncooperative games in specific cases and applications.

Stochastic Differential Games

This part is devoted to the description of Stochastic differential games which are games with randomness in the dynamics.

The dynamics is given by a stochastic differential equation

$$dX_t = f(X_t, u_t, v_t)dt + \sigma(X_t, u_t, v_t)dW_t, \quad (10)$$

where W is a Brownian motion on a given probability space. Victor and Ursula choose the controls u and v which are adapted processes (with respect to the filtration generated by the Brownian motion). The notions of nonanticipative strategies are similar to the determinist case but the strategies are nonanticipative in time *almost surely*. We refer the reader to Cardaliaguet and Rainer (2013) for a precise formulation of strategies. The rigorous description of the game is rather technical due to the need of stochastic analysis. The reader can refer to the nice article (Rainer 2007) for a precise formulation. We want just to stress some important aspects of stochastic zero-sum games. Of course the notions of qualitative and quantitative games are relevant. The qualitative aspect is not yet well-studied. Oppositely the quantitative aspect is now well-studied. Consider for instance a payoff of the form

$$C(t_0, x_0, u, v) := E \left[\int_{t_0}^T L(x(s), u(s), v(s))ds + g(x(T)) \right],$$

where E is the expectation. The existence of the value was obtained in Fleming and Souganidis (1989), by adopting the same scheme that which was described in section “Quantitative Target Games”: The values are the unique viscosity solution of a – second order – Hamilton Jacobi Isaacs Equation which has a unique solution. Another approach restricted to the case where the drift σ of Eq. (10) is not degenerate is possible with a

different notion of strategy and by using backward stochastic differential equations (Hamadene and Lepeltier 1995).

The nonantagonist case could be solved similarly than the zero-sum case up to hard technicalities due to stochastic analysis (Buckdahn et al. 2004).

Differential Games with Incomplete Information

This concerns the case where the players have not the same information during the game. This was studied in the case of (discrete) repeated games since the pioneering work (Aumann and Maschler 1995) with the help of mixed strategies (Petrosjan 2004).

We describe now a zero-sum differential game with incomplete information on the initial position. Consider the two players quantitative game of section “Qualitative and Quantitative Differential Games,” with dynamics (3) and cost (8). We suppose that the initial position belongs to a set of I possible initial positions

$$x_0^i, i = 1, \dots, I.$$

Let us describe how the game is played.

$$\left\{ \begin{array}{l} \min \left\{ \frac{\partial V}{\partial t}(t, X, p) + H \left(x, \frac{\partial V}{\partial t}(t, X, p) \right); \lambda_{\min} \left(\frac{\partial^2 V}{\partial p^2}(t, X, p) \right) \right\} = 0, \text{ for all } (t, X, p) \in [0, T[\times \mathbb{R}^{N_I} \times \Delta(I) \\ V(T, X) = \sum_i p_i g(x_i), \text{ for all } (x, p) \in \mathbb{R}^{N_I} \times \Delta(I) \end{array} \right. \quad (11)$$

where $X = (x_i)_{i=1,2,\dots,I}$, with the Hamiltonian

$$H(X, \xi) := \min_{v \in V} \max_{u \in U} \sum_{i=1}^I \{ L(x_i, u, v) + \langle f(x_i, u, v), \xi_i \rangle \}$$

and where $\lambda_{\min} \left(\frac{\partial^2 V}{\partial p^2}(t, X, p) \right)$ is the smallest eigenvalue of the symmetric matrix $\frac{\partial^2 V}{\partial p^2}(t, X, p)$. This shows the existence of a Value in mixed non-anticipative strategies (Cardaliaguet 2007). This

Before the game starts, the integer i is chosen randomly according to a probability p (belonging to the set $\Delta(I)$ of probabilities on $\{1, 2, \dots, I\}$). The index i is communicated to Ursula only and not to Victor.

As previously Victor wants to minimize the cost C and Ursula want to maximize it. Both players observe – in a nonanticipative way – the actions (the controls) played by his/her opponent. It is worth pointing out that Victor has not enough information to be able compute the current position of the game. Nevertheless by observing his opponents behavior, he can try to deduce from this observation his missing information. Moreover knowing this, Victor will try to play in such way that he can hide as much as is possible his information in order to keep an advantage (the way of hiding the information is the choice of random strategies).

This allows to define values which depends on (t, x) but also on the probability p : $V^\#(t, x, p)$ and $V^b(t, x, p)$. Under a suitable Isaacs condition, both value coincide because there are the unique viscosity solution of the following second order Hamilton Jacobi Isaacs equation

model was also extended to both side incomplete information: Namely, the initial position is randomly chosen among $x_0^{ij}, i = 1, 2, \dots, I, j = 1, \dots, J$ according to a probability $p \otimes q \in \Delta(I) \times \Delta(J)$, the index i is communicated to Ursula but not to Victor while the index j is communicated only to Victor. Also the information can concern not only the initial positions $x_0^{i,j}$ but also the dynamics $f_{i,j}$ or the costs. The existence of the Value is obtained in these cases and also for stochastic differential games (Cardaliaguet and Rainer 2009).

Observe that the incomplete information is of a finite nature, one can consider the case where the initial position is chosen randomly according to a probability measure μ_0 not necessarily finite. The above model concerns the case where $\mu_0 = \sum_{i=1}^I \delta_{x_i}$. Then it is still possible to obtain the existence of the value and moreover to obtain the existence of the value in pure strategies for initial measure which are absolutely continuous (Cardaliaguet et al. 2014). However, the Hamilton Jacobi Isaacs equation have to be stated on the space of probability measures (Cardaliaguet et al. 2014), this is technically difficult and exceed the scope of the present Note. Let us just mention that this allows to study differential games with measured dynamics (Marigonda and Quincampoix 2018).

In the model above, the incomplete information is a data before the game starts; there is no further change in the information. A partial revealing during the game would be a very interesting model, unfortunately very few is known in this context and the subject is still widely open.

Miscellaneous

Several fast growing areas concerns differential games. Among them we just mention mean field games. This could be viewed as non zero sum differential games with a very large number of Players. The number of players is so large that an individual player has almost no influence on the game but the game evolves according the “mean” of the actions of the Players. This theory was developed by Lasry and Lions (cf. for instance (Lasry and Lions 2007)). Mean field games are the limit of Nash differential games when the number of the players tends to infinity, so the model involves both the limit $m(t)$ – as $N \rightarrow +\infty$ – of the position $\frac{1}{N} \sum_{i=1}^N \delta_{x_i(t)}$ of each player $x_i(t)$ and a value function satisfying an Hamilton Jacobi Isaacs equation depending on $m(t)$. We refer the reader to Buckdahn et al. (2009b), Buckdahn et al. (2009a), and Cardaliaguet (2017) for more details.

Bibliography

Primary Literature

- Aubin J-P (1991) Viability theory. Birkhäuser, Boston
- Aumann RJ, Maschler M (1995) Repeated games with incomplete information. MIT Press, Cambridge
- Bardi M, Capuzzo Dolcetta I (1997) Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations. Systems and control: foundations and applications. Birkhäuser, Boston. xvii, 570 pp
- Basar T, Bernhard P (1995) H^∞ -optimal control and related minimax design problems. In: A dynamic game approach, 2nd edn. Systems and control: foundations and applications. Birkhäuser, Boston
- Berkovitz LD (1994) A theory of differential games. In: Basar T et al (ed) Advances in dynamic games and applications. Birkhäuser, Boston. Ann Int Soc Dyn Games 1:3–22
- Bernhard P (1988) Differential games. In: Singh MG (ed) Systems and control encyclopedia, theory technology application. Pergamon Press, Oxford
- Bettiol P, Cardaliaguet P, Quincampoix M (2006) Zero-sum state constrained differential games: existence of value for Bolza problem. Int J Game Theory 34(4):495–527
- Breakwell JV (1977) Zero-sum differential games with terminal payoff. In: Hagedorn P, Knobloch HW, Olsder GH (eds) Differential game and applications. Lecture notes in control and information sciences, vol 3. Springer, Berlin
- Bressan A (2011) Noncooperative differential games. Milan J Math 79:357–427
- Buckdahn R, Cardaliaguet P, Rainer C (2004) Nash equilibrium payoffs for nonzero-sum stochastic differential games. SIAM J Control Optim 43(2):624–642
- Buckdahn R, Djehiche B, Li J, Peng S (2009a) Mean-field backward stochastic differential equations: a limit approach. Ann Probab 37:1524–1565
- Buckdahn R, Li J, Peng S (2009b) Mean-field backward stochastic differential equations and related partial differential equations. Stoch Process Appl 119:3133–3154
- Buckdahn R, Li J, Quincampoix M (2013) Value function of differential games without Isaacs conditions. An approach with nonanticipative mixed strategies. Int J Game Theory 42(4):989–1020
- Buckdahn R, Li J, Quincampoix M (2014) Value in mixed strategies for zero-sum stochastic differential games without Isaacs condition. Ann Probab 42(4):1724–1768
- Buckdahn R, Quincampoix M, Rainer C, Xu Y (2016) Differential games with asymmetric information and without Isaacs’ condition. Int J Game Theory 45(4):795–816
- Cardaliaguet P (1996) A differential game with two players and one target. SIAM J Control Optim 34(4):1441–1460
- Cardaliaguet P (1997) Nonsmooth semi-permeable barriers, Isaacs’ equation, and application to a differential game with one target and two players. Appl Math Optim 36:125–146

- Cardaliaguet P (2007) Differential games with asymmetric information. *SIAM J Control Optim* 46(3):816–838
- Cardaliaguet P (2017) The convergence problem in mean field games with local coupling. *Appl Math Optim* 76(1):177–215
- Cardaliaguet P, Plasckacz S (2003) Existence and uniqueness of a Nash equilibrium feedback for a simple non zero-sum differential game. *Int J Game Theory* 32(4):561–593
- Cardaliaguet P, Quincampoix M (2008) Determinist differential games under probability knowledge of initial condition. *Int Game Theory Rev* 10(1):1–16
- Cardaliaguet P, Rainer C (2009) Cardaliaguet, Pierre; Rainer, Catherine stochastic differential games with asymmetric information. *Appl Math Optim* 59(1):136
- Cardaliaguet P, Rainer C (2013) Pathwise strategies for stochastic differential games. *Appl Math Optim* 68(1):75–84
- Cardaliaguet P, Quincampoix M, Saint-Pierre P (1999) Set-valued numerical analysis for optimal control and differential games. In: Bardi M, Parthasarathy T, TES R (eds) *Numerical methods for optimal control and numerical games*. Annals of International Society of Dynamical Games. Birkhäuser, Boston, pp 177–249
- Cardaliaguet P, Quincampoix M, Saint-Pierre P (2001) Pursuit differential games with state constraints. *SIAM J Control Optim* 39(5):1615–1632
- Cardaliaguet P, Jimenez C, Quincampoix M (2014) Pure and random strategies in differential game with incomplete informations. *J Dyn Games* 1(3):363–375
- Elliot N, Kalton N (1972) The existence of value in differential games. *Mems Am Math Soc* 126
- Evans LC, Souganidis PE (1984) Differential games and representation formulas for solutions of Hamilton-Jacobi equations. *Indiana Univ Math J* 282:487–502
- Fleming W, Souganidis P (1989) On the existence of value functions of two-player, zero-sum stochastic differential games. *Indiana Univ Math J* 38(2):293–314
- Flynn J (1973) Lion and man: the boundary constraints. *SIAM J Control* 11:397
- Hamadene S, Lepeltier J-P (1995) Backward equations, stochastic control and zero-sum stochastic differential games. *Stoch Int J Probab Stoch Process* 54:221–231
- Jimenez C, Quincampoix M (2018) Hamilton Jacobi Isaacs equations for differential games with asymmetric information on probabilistic initial condition. *J Math Anal Appl* 457(2):1422–1451
- Kleimenov AF (1993) Nonantagonistic positional differential games. Nauka, Ekaterinburg
- Lasry J-M, Lions P-L (2007) Mean field games. *Jpn J Math* 2(1):229–260
- Marigonda A, Quincampoix M (2018) Mayer control problem with probabilistic uncertainty on initial positions. *J Differ Equ* 264(5):3212–3252
- Petrosjan LA (2004) Cooperation in games with incomplete information. In: *Nonlinear analysis and convex analysis*. Yokohama Publishers, Yokohama, pp 469–479
- Plaskacz S, Quincampoix M (2000) Discontinuous Mayer control problem under state-constraints. *Topol Methods Nonlinear Anal* 15:91–100
- Pontryagin N (1968) Linear differential games, I–II. *Soviet Math Dokl* 8(3–4):769–771, 910–913
- Quincampoix M (1992) Differential inclusions and target problems. *SIAM J Control Optim* 30(2):324–335
- Quincampoix M, Veliov V (2005) Optimal control of uncertain systems with incomplete information for the disturbance. *SIAM J Control Optim* 43(4):1373–1399
- Rainer C (2007) On two different approaches to nonzero-sum stochastic differential games. *Appl Math Optim* 56:131–144
- Roxin E (1969) The axiomatic approach in differential games. *J Optim Theory Appl* 3:153–163
- Subbotin AI (1995) Generalized solutions of first-order PDEs. The dynamical optimization perspective. Translated from the Russian. *Systems & control: foundations & applications*. Birkhäuser, Boston
- Varaiya P (1967) The existence of solution to a differential game. *SIAM J Control Optim* 5:153–162
- Varaiya P, Lin J (1967) Existence of saddle points in differential game. *SIAM J Control Optim* 7(1):141–157
- Von Neumann J, Morgenstern O (1946) *Theory of games and economic behaviour*. Princeton University Press, Princeton

Books and Reviews

- Bardi M, Ragghavan TES, Parthasarath T (eds) (1999) *Stochastic and differential games. Theory and numerical methods*. Dedicated to Prof. A. I. Subbotin. Annals of the International Society of Dynamic Games, vol 4. Birkhäuser, Boston
- Basar T, Olsder GJ (1999) *Dynamic noncooperative game theory*, 2nd edn. Classics in applied mathematics, vol 23. SIAM, Society for Industrial and Applied Mathematics, Philadelphia
- Blaquière A, Gérard F, Leitman G (1969) *Quantitative and qualitative games*. Academic, New York
- Buckdahn R, Cardaliaguet P, Quincampoix M (2011) Some recent aspects of differential game theory. *Dyn Games Appl* 1(1):74–114
- Dockner E, Jörgensen S, Van Long N, Sorger G (2000) *Differential games in economics and management science*. Cambridge University Press, Cambridge
- Hajek O (1975) *Pursuit games*. Academic, New York
- Isaacs R (1965) *Differential games*. Wiley, New York
- Jorgensen S, Quincampoix M, Vincent T (eds) (2007) *Advances in dynamic games theory*. Annals of International Society of Dynamic Games. Birkhäuser, Boston
- Krasovski NN, Subbotin AI (1988) *Game-theoretical control problems*. Springer, New York
- Melikyan AA (1998) *Generalized characteristics of first order PDEs. Applications in optimal control and differential games*. Birkhäuser, Boston
- Patsko VS, Turova VL (2000) *Numerical study of differential games with the homicidal chauffeur dynamics*. Russian Academy of Sciences, Institute of Mathematics and Mechanics, Ekaterinburg
- Petrosyan LA (1993) *Differential games of pursuit*. Series on optimization, vol 2. World Scientific, Singapore

Mechanism Design

Ron Lavi

The Technion – Israel Institute of Technology,
Haifa, Israel

Article Outline

Glossary

Definition of the Subject

Introduction

Formal Model and Early Results

Quasi-Linear Utilities and the VCG Mechanism

The Importance of the Domain's Dimensionality

Single-Dimensional Domains

Multi-dimensional Domains

Budget Balancedness and Bayesian Mechanism
Design

Interdependent Valuations

Future Directions

Bibliography

Glossary

A social choice function A function that determines a social choice according to players' preferences over the different possible alternatives.

A mechanism A game in incomplete information, in which player strategies are based on their private preferences. A mechanism implements a social choice function f if the equilibrium strategies yield an outcome that coincides with f .

Dominant strategies An equilibrium concept where the strategy of each player maximizes her utility, no matter what strategies the other players choose.

Bayesian-Nash equilibrium An equilibrium concept that requires the strategy of each player to maximize the expected utility of the player, where the expectation is taken over the types of the other players.

VCG mechanisms A family of mechanisms that implement in dominant strategies the social choice function that maximizes the social welfare.

Definition of the Subject

Mechanism design is a sub-field of economics and game theory that studies the construction of social mechanisms in the presence of rational but selfish individuals (players/agents). The nature of the players dictates a basic contrast between the social planner, that aims to reach a socially desirable outcome, and the players, that care only about their own private utility. The underlying question is how to incentivize the players to cooperate, in order to reach the desirable social outcomes. A *mechanism* is a game, in which each agent is required to choose one action among a set of possible actions. The social designer then chooses an *outcome*, based on the chosen actions. This outcome is typically a coupling of a physical outcome, and a payment given to each individual. Mechanism design studies how to design the mechanism such that the *equilibrium* behavior of the players will lead to the socially desired goal. The theory of mechanism design has greatly influenced several sub-fields of micro-economics, for example auction theory and contract theory, and the 2007 Nobel prize in Economics was awarded to Leonid Hurwicz, Eric Maskin, and Roger Myerson “for having laid the foundations of mechanism design theory”.

Introduction

It will be useful to start with an example of a mechanism design setting, the well-known “public project” problem (Clarke [8]): a government is trying to decide on a certain public project (the common example is “building a bridge”). The

project costs C dollars, and each player, i , will benefit from it to an amount of v_i dollars, where this number is known only to the player herself. The government desires to build the bridge if and only if $\sum_i v_i > C$. But how should this condition be checked? Clearly, every player has an interest in over-stating its own v_i , if this report is not accompanied by any payment at all, and most probably agents will understate their values, if asked to pay some proportional amount to the declared value. Clarke describes an elegant mechanism that solves this problem. His mechanism has the fantastic property that, from the point of view of every player, no matter what the other players declare, it is always *in the best interest* of the player to declare his true value. Thus, truthful reporting is a dominant-strategy equilibrium of the mechanism, and under this equilibrium, the government's goal is fully achieved. A more formal treatment of this result is given in section “[Quasi-Linear Utilities and the VCG Mechanism](#)” below.

Clarke's paper, published in the early 1970s, was part of a large body of work that started to investigate mechanism design questions. Most of the early works used two different assumptions about the structure of players' utilities. Under the assumption that utilities are general, and that the influence of monetary transfers on the utility are not well predicted, the literature have produced mainly impossibilities, which are described in section “[Formal Model and Early Results](#)”. The assumption that utilities are quasi-linear in money was successfully used to introduce positive and impressive results, as discussed in detail in sections “[Quasi-Linear Utilities and the VCG Mechanism](#)” and “[The Importance of the Domain's Dimensionality](#)”. These mechanisms apply the solution concept of dominant-strategy equilibrium, which is a strong solution concept that may prevent several desirable properties from being achieved. To overcome its difficulties, the weaker concept of a Bayesian–Nash equilibrium is usually employed. This concept, and one main possibility result that it provides, are described in section “[Budget Balancedness and Bayesian Mechanism Design](#)”. The last important model that this entry covers aims to capture settings where the players' values

are not fully observed by each player separately. Rather, each player receives a *signal* that gives a partial indication to her valuation. Mechanism design for such settings is discussed in section “[Interdependent Valuations](#)”.

One of the most impressive applications of the general mechanism design literature is auction theory. An auction is a specific form of a mechanism, where the outcome is simply the specific allocation of the goods to the players, plus the prices they are required to pay. Vickrey [27] initiated the study of auctions in a mechanism design setting, and in fact perhaps the study of mechanisms itself. After the fundamental study general mechanism design in the 1970s, in the 1980s the focus of the research community returned to this important application, and many models were studied.

We note that there are several other entries in this book that are strongly related to the subject of “mechanism design”. In particular, the entry on “[► Game Theory, Introduction to](#)” gives a broader background on the mathematical methods and tools that are used by mechanism designers, and the entry on “[► Implementation Theory](#)” handles similar subjects to this entry from a different point of view.

Formal Model and Early Results

A social designer wishes to choose one possible outcome/alternative out of a set A of possible alternatives. There are n players, each has her own preference order $\succeq_i$ over A . This preference order is termed the player's “type”. The set (domain) of all valid player preferences is denoted by V_i . The designer has a social choice function $f: V_1 \times \cdots \times V_n \rightarrow A$, that specifies the desired alternative, given any profile of individual preferences over the alternatives. The problem is that these preferences are private information of each player – the social designer does not know them, and thus cannot simply invoke f in order to determine the social alternative. Players are assumed to be strategic, and therefore we are in a game-theoretic situation.

To *implement* the social choice function, the designer constructs a “game in incomplete

information”, as follows. Each player is required to choose an action out of a set of possible actions $\mathcal{A}_i$, and a target function $g : \mathcal{A}_1 \times \cdots \times \mathcal{A}_n \rightarrow A$ specifies the chosen alternative, as a function of the players’ actions. A player’s choice of action may, of-course, depend on her actual preference order. Furthermore, we assume an incomplete information setting, and therefore it cannot depend on any of the other players’ preferences. Thus, to play the game, player i chooses a strategy $s_i : V_i \rightarrow \mathcal{A}_i$.

A strategy $s_i(\cdot)$ dominates another strategy $s'_i(\cdot)$ if, for every tuple of actions a_{-i} of the other players, and for every preference $\succeq_i \in V_i$, $g(s_i(\succeq_i), a_{-i}) \succeq_i g(s'_i(\succeq_i), a_{-i})$, for any $a_i \in \mathcal{A}_i$. In other words, no matter what the other players are doing, the player cannot improve her situation by using an action other than $s_i(\succeq_i)$.

A mechanism *implements* the social choice function f in dominant strategies if there exist dominant strategies $s_1(\cdot), \dots, s_n(\cdot)$ such that $f(\succeq_1, \dots, \succeq_n) = g(s_1(\succeq_1), \dots, s_n(\succeq_n))$, for any profile of preferences $\succeq_1, \dots, \succeq_n$. In other words, a mechanism implements the social choice function f if, given that players indeed play their equilibrium strategies (in this case the dominant strategies equilibrium), the outcome of the mechanism coincides with f ’s choice.

The theory of mechanism design asks: given a specific problem domain (an alternative set and a domain of preferences), and a social choice function, how can we construct a mechanism that implements it (if at all)? As we shall see below, the literature uses a variety of “solution concepts”, in addition to the concept of dominant strategies equilibrium, and an impressive set of understandings have emerged.

The concept of implementing a function with a dominant-strategy mechanism seems at first too strong, as it requires each player to know exactly what action to take, regardless of the actions the others take. Indeed, as we will next describe in detail, if we do not make any further assumptions then this notion yields mainly impossibilities. Nevertheless, it is not completely empty, and it may be useful to start with a positive example, to illustrate the new notions defined above.

Consider a voting scenario, where the society needs to choose one out of two candidates. Thus,

the alternative set contains two alternatives (candidate 1” and candidate 2”), and each player either prefers 1 over 2, 2 over 1, or is indifferent between the two. It turns out that the majority voting rule is the dominant strategy implementable, by the following mechanism: each player reports her top candidate, and the candidate that is preferred by the majority of the players is chosen. This mechanism is a “direct-revelation” mechanism, in the sense that the action space of each player is to report a preference, and g is exactly f . In a direct-revelation mechanism, the hope is that truthful reporting (i.e. $s_i(\succeq_i) = \succeq_i$) is a dominant strategy. It is not hard to verify that in this two candidates setting, this is indeed the case, and hence the mechanism implements in dominant-strategies the majority voting rule.

An elegant generalization for the case of a “single-peaked” domain is as follows. Assume that the alternatives are numbered as $A = \{a_1, \dots, a_n\}$, and the valid preferences of a player are single-peaked, in the sense that the preference order is completely determined by the choice of a peak alternative, a_p . Given the peak, the preference between any two alternatives a_i, a_j is determined according to their distance from a_p , i.e. $a_i \succeq_i a_j$ if and only if $|j - p| \leq |i - p|$. Now consider the social choice function $f(p_1, \dots, p_n) = \text{median}(p_1, \dots, p_n)$, i.e. the chosen alternative is the median alternative of all peak alternatives.

Theorem 1 *Suppose that the domain of preferences is single-peaked. Then the median social choice function is implementable in dominant strategies.*

Proof Consider the direct revelation mechanism in which each player reports a peak alternative, and the mechanism outputs the median of all peaks. Let us argue that reporting the true peak alternative is a dominant strategy. Suppose the other players reported p_{-i} , and that the true peak of player i is p_i . Let p_m be the median index. If $p_i = p_m$ then clearly player i cannot gain by declaring a different peak. Thus, assume that $p_i < p_m$, and let us examine a false declaration p'_i player i . If $p'_i \leq p_m$ then p_m remains the median, and the player did not gain. If $p'_i > p_m$ then the

new median is $p'_m \geq p_m$, and since $p_i < p_m$, this is less preferred by i . Thus, player i cannot gain by declaring a false peak alternative if the true peak alternative is smaller or equal to the median alternative. A similar argument holds for the case of $p_i > p_m$. $\square$

In a voting situation with two candidates, the median rule becomes the same as the majority rule, and the domain is indeed single-peaked. When we have three or more candidates, it is not hard to verify that the majority rule is different than the median rule. In addition, one can also check that the direct-revelation mechanism that uses the majority rule does not have truthfulness as a dominant strategy. Of-course, many times one cannot order the candidates on a line, and any preference ordering over the candidates is plausible. What voting rules are implementable in such a setting? This question was asked by Gibbard [12] and Satterthwaite [26], who provided a beautiful and fundamental impossibility. A domain of player preferences is unrestricted if it contains all possible preference orderings. In our voting example, for instance, the domain is unrestricted if every ordering of the candidates is valid (in contrast to the case of a single-peaked domain). A social choice function is dictatorial if it always chooses the top alternative of a certain fixed player (the dictator).

Theorem 2 ([12, 27]) *Every social choice function over an unrestricted domain preferences, with at least three alternatives, must be dictatorial.*

The proof this theorem, and in fact of most other impossibility theorems in mechanism design, uses as a first step the powerful direct-revelation principle. Though the examples we have seen above use a direct revelation mechanism, one can try to construct “complicated” mechanisms with “crazy” action spaces and outcome functions, and by this obtain dominant strategies. How should one reason about such vast space of possible constructions? The revelation principle says that one cannot gain extra power by such complex constructions, since if there exists an implementation to a specific function then there exists a direct revelation mechanism that implements it.

Theorem 3 (The Direct Revelation Principle)

Any implementable social choice function can also be implemented (using the same solution concept) by a direct-revelation mechanism.

Proof Given a mechanism M that implements f , with dominant strategies $s_i^*(\cdot)$, we construct a direct revelation mechanism M' as follows: for any tuple of preferences $\succeq = (\succeq_1, \dots, \succeq_n)$, $g'(\succeq) = g(s^*(\succeq))$. Since $s_i^*(\cdot)$ is a dominant strategy in M , we have that for any fixed $\succeq_{-i} \in V_{-i}$ and any $\succeq_i \in V_i$, the action $a_i = s_i^*(\succeq_i)$ is dominant when i 's type is $\succeq_i$. Hence declaring any other type $\tilde{\succeq}_i$ that will “produce” an action $\tilde{a}_i = s_i^*(\tilde{\succeq}_i)$, cannot increase i 's utility. Therefore, the strategy $\succeq_i$ in M' is dominant. $\square$

The proof uses the dominant-strategies solution concept, but any other equilibrium definition will also work, using the same argumentation. Though technically very simple, the revelation principle is fundamental. It states that, when checking if a certain function is implementable, it is enough to check the direct-revelation mechanism that is associated with it. If it turns out to be truthful, we still may want to implement it with an indirect mechanism that will seem more natural and “real”, but if the direct-revelation mechanism is not truthful, then there is no hope of implementing the function.

The proof of the theorem of Gibbard and Satterthwaite relies on the revelation principle to focus on direct-revelation mechanisms, but this is just the beginning. The next step is to show that *any* non-dictatorial function is non-implementable. The proof achieves this by an interesting reduction to Arrow's theorem, from the field of social choice theory. This theory is concerned with the possibilities and impossibilities of social preference aggregations that will exhibit desirable properties. A social welfare function $F : V \rightarrow \mathcal{R}$ aggregates the individuals' preferences into a single preference order over all alternatives, where $\mathcal{R}$ is the set of all possible preference orders over A . Arrow [2] describes few desirable properties from a social welfare function, and shows that no social choice function can satisfy all:

Definition 1 (Arrow's desirable properties)

1. A social welfare function satisfies “weak

- Pareto” if whenever all individuals strictly prefer alternative a to alternative b then, in the social preference, a is strictly preferred to b .
2. A social welfare function is “a dictatorship” if there exists an individual for which the social preference is always identical to his own preference.
 3. A social welfare function F satisfies the “Independence of Irrelevant Alternatives” property (IIA) if, for any preference orders $R, \tilde{R} \in \mathcal{R}$ and any $a, b \in A$,

$$a \succ_{F(R)} b \text{ and } b \succ_{F(\tilde{R})} a \Rightarrow \exists i : a \succ_{R_i} b \text{ and } b \succ_{R_i} a$$

(where $a \succ_{R_i} b$ iff a is preferred over b in R_i). In other words, if the social preference between a and b was flipped when the individual preferences were changed from R to $\tilde{R}$, then it must be the case that some individual flipped his own preference between a and b .

Arrow’s impossibility theorem holds for the *unrestricted* domain of preferences, i.e. when all preference orders are possible:

Theorem 4 ([2]) *Assume $|A| \geq 3$. Any social welfare function over an unrestricted domain of preferences that satisfy both weak Pareto and Independence of Irrelevant Alternatives must be a dictatorship.*

Gibbard and Satterthwaite’s proof reveals an interesting and important connection between the concept of implementation in dominant strategies, and Arrow’s condition IIA. The proof shows how to construct, from a given implementable social choice function f , a social welfare function, F , that satisfies IIA and weak Pareto. In addition, F always places the alternative chosen by f as the most preferred alternative. By Arrow’s theorem, the resulting social welfare function must be dictatorial. In turn, this implies that f is dictatorial. The construction of F from f is the straightforward one: the top alternative is f ’s choice to the original preferences, say a . Then a is lowered to be the least preferred alternative in all the preferences, and f ’s new choice is placed second,

etc. The interesting exercise is to show that the implementability of F implies that F satisfies Arrow’s conditions. In fact, as the proof shows that any implementable social choice function f entails a social welfare function F that “extends” f and satisfies Arrow conditions, it actually provides a strong argument for the reasonability of Arrow’s requirement—they are simply implied by the implementability requirement.

In view of these strong impossibility results, it is natural to ask whether the entire concept of a mechanism can yield positive constructions. The answer is a big yes, under the “right” set of assumptions, as discussed in the next sections.

Quasi-Linear Utilities and the VCG Mechanism

The model formalization of the previous section ignores the existence of money, or, more accurately, the fact that it has a more or less predictable effect on a player’s utility. The quasi-linear utilities model takes this into account, and players are assumed to have monetary value for each alternative.

Formally, the type of a player is a valuation function $v_i: A \rightarrow \mathfrak{R}$ that describes the monetary value that the player will obtain from each chosen alternative (as before v_i is taken from a domain of valid valuations V_i and $V = V_1 \times \cdots \times V_n$). Note that the value of a player does not depend on the other players’ values (this is termed the private value assumption). The mechanism designer can now additionally pay each player (or charge money from her), and the total utility of player i if the chosen outcome is a and in addition she pays a price P_i is $v_i(a) - P_i$. A direct mechanism for quasi-linear utilities includes an outcome function $f: V \rightarrow A$ (as before), as well as price functions $p_i: V \rightarrow \mathfrak{R}$ for each player i (the definition of an indirect mechanism is the natural parallel of the definition of the previous section; the revelation principle holds for quasi-linear utilities as well, and we focus here on direct mechanisms). The implicit assumption is that a player aims to maximize her resulting utility, $v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i})$, and this leads us to the definition of a truthful mechanism, that parallels that of the previous section:

Definition 2 (Truthfulness, or Incentive Compatibility, or Strategy-Proofness) A direct revelation mechanism is “truthful” (or incentive-compatible, or strategy-proof) if the dominant strategy of each player is to reveal her true type, i.e. if for every $v_{-i} \in V_{-i}$ and every $v_i, v'_i \in V_i$,

$$v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i}) \geq v_i(f(v'_i, v_{-i})) - p_i(v'_i, v_{-i})$$

Using this framework, we can return to the example from section “[Introduction](#)” (“building a bridge”), and construct a truthful mechanism to solve it. Recall that, in this problem, a government is trying to decide on a certain public project, which costs C dollars. Each player, i , will benefit from it to an amount of v_i dollars, where this number is known only to the player herself. The government desires to build the bridge if and only if $\sum_i v_i \geq C$. Clarke [8] designed the following mechanism. Each player reports a value, $\tilde{v}_i$, and the bridge is built if and only if $\sum_i \tilde{v}_i \geq C$. If the bridge is not built, the price of each player is 0. If the bridge is built then each player, i , pays the minimal value she could have declared to maintain the positive decision. More precisely, if $\sum_{i' \neq i} \tilde{v}_{i'} \geq C$ then she still pays zero, and otherwise she pays $C - \sum_{i' \neq i} \tilde{v}_{i'}$.

Theorem 5 *Bidding the true value is a dominant strategy in the Clarke mechanism.*

Proof Consider the truthful bidding for player i , v_i , vs. another possible bid $\tilde{v}_i$ (fixing the bids of the other players to arbitrarily be $\tilde{v}_{-i}$). If with v_i the project was rejected then $v_i < C - \sum_{i' \neq i} \tilde{v}_{i'}$. In order to change the decision to an accept, the player would need to declare $\tilde{v}_i \geq C - \sum_{i' \neq i} \tilde{v}_{i'}$. In this case i 's payment will be $C - \sum_{i' \neq i} \tilde{v}_{i'}$ which is smaller than v_i , as observed above. Thus, i 's resulting utility will be negative, hence bidding $\tilde{v}_i$ did not improve her utility.

On the other hand, assume that with v_i the project is accepted. Therefore, the player's utility from declaring v_i is non-negative. Note that the price that the player pays in case of an accept does not depend on her bid. Thus, the only way to

change i 's utility (if at all) is to declare some $\tilde{v}_i$ that will cause the project to be rejected. But in this case i 's utility will be zero, hence she did not gain any benefit. $\square$

Subsequently, Groves [13] made the remarkable observation that Clarke's mechanism is in fact a special case of a much more general mechanism, that solves the welfare maximization problem on *any* domain with private values and quasi-linear utilities. For a given set of player types $v_1, \dots, v_n$, the *welfare* obtained by an alternative $a \in A$ is $\sum_i v_i(a)$. A social choice function is termed a *welfare maximizer* if $f(v)$ is an alternative with maximal welfare, i.e. $f(v) \in \operatorname{argmax}_{a \in A} \{\sum_{i=1}^n v_i(a)\}$.

Definition 3 (VCG Mechanisms) Given a set of alternatives A , and a domain of players' types $V = V_1 \times \dots \times V_n$, a VCG mechanism is a direct revelation mechanism such that, for any $v \in V$,

1. $f(v) \in \operatorname{argmax}_{a \in A} \{\sum_{i=1}^n v_i(a)\}$.
2. $p_i(v) = -\sum_{j \neq i} v_j(f(v)) + h_i(v_{-i})$, where $h_i: V_{-i} \rightarrow \Re$ is an arbitrary function.

Ignore for a moment the term $h_i(v_{-i})$ in the payment functions. Then the VCG mechanism has a very natural interpretation: it chooses an alternative with maximal welfare according to the reported types, and then, by making additional payments, it equates the utility of each player to that maximal welfare level.

Theorem 6 ([13]) *Any VCG mechanism truthfully implements the welfare maximizing social choice function.*

Proof We argue that $s_i(v_i) = v_i$ is a dominant strategy for i . Fix any $v_{-i} \in V_{-i}$ as the declarations (actions) of the other players, any $v'_i \neq v_i$, and assume by contradiction that $v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i}) < v_i(f(v'_i, v_{-i})) - p_i(v'_i, v_{-i})$. Replacing $p_i(\cdot)$ with the specific VCG payment function, and eliminating the term $h_i(v_{-i})$ from both sides, we get: $v_i(f(v_i, v_{-i})) + \sum_{j \neq i} v_j(f(v_i, v_{-i})) < v_i(f(v'_i, v_{-i})) + \sum_{j \neq i} v_j(f(v'_i, v_{-i}))$. Therefore, it must be that $f(v_i, v_{-i}) \neq f(v'_i, v_{-i})$. Denote

$f(v_i, v_{-i}) = a$ and $f(v'_i, v_{-i}) = b$. The above equation is now $v_i(a) + \sum_{j \neq i} v_j(a) < v_i(b) + \sum_{j \neq i} v_j(b)$, or, equivalently, $\sum_{i=1}^n v_i(a) < \sum_{i=1}^n v_i(b)$, a contradiction to the fact that $f(v_i, v_{-i}) = a$, since $f(\cdot)$ is a welfare maximizer.

□

Thus, we see that the welfare maximizing social choice function can always be implemented, no matter what the problem domain is, under the assumption of quasi-linear utilities. The VCG mechanism is named after Vickrey, whose seminal paper [27] on auction theory was the first to describe a special case of the above mechanism (this is the second price auction; see the entry on auction theory for more details), after Clarke, who provided the second example, and after Groves himself, that finally pinned down the general idea.

Clarke's work can be viewed, in retrospect, as a suggestion for one specific form of the function $h_i(v_{-i})$, namely $h_i(v_{-i}) = \sum_{j \neq i} v_j(f(v_{-i}))$ (this is a slight abuse of notation, as f is defined for n players, but the intention is the straight-forward one — f chooses an alternative with maximal welfare). This form for the $h_i(\cdot)$'s gives the following property: if a player does not influence the social choice, her payment is zero, and, in general, a player pays the “monetary damage” to the other players (i.e. the welfare that the others lost) as a result of i 's participation. Additionally, with Clarke's payments, a truthful player is guaranteed a non-negative utility, no matter what the others declare. This last property is termed “individual rationality”.

The Importance of the Domain's Dimensionality

The impressive property of the VCG mechanism is its generality with respect to the domain of preferences — it can be used for *any* domain. On the other hand, VCG is restrictive in the sense that it can be used only to implement one specific goal, namely welfare maximization. Given the possibility that VCG presents, it is natural to ask if the assumption of quasi-linear utilities and private values allows the designer to implement many

other different goals. It turns out that the answer depends on the “dimensionality” of the domain, as is discussed in this section.

Single-Dimensional Domains

Consider first a domain of preferences for which the type $v_i(\cdot)$ can be completely described by a single number v_i , in the following way. For each player i , a subset of the alternatives are “losing” alternatives, and her value for all these alternatives is always 0. The other alternatives are “winning” alternatives, and the value for each “winning” alternative is the same, regardless of the specific alternative. Such a domain is “single dimensional” in the sense that one single number completely describes the entire valuation vector. As before, this single number (the value for winning), is private to the player, and here this is the *only* private information of the player. The public project domain discussed above is an example of a single-dimensional domain: the losing alternative is the rejection of the project, and the winning alternative is the acceptance of the project.

A major drawback of the VCG mechanism, in general, and with respect to the public project domain in particular, is the fact that the sum of payments is not balanced (a broader discussion on this is given in section “[Budget Balancedness and Bayesian Mechanism Design](#)” below). In particular, the payments for the public project domain may not cover the entire cost of the project. Is there a different mechanism that always covers the entire cost? The positive answer that we shall soon see crucially depends on the fact that the domain is single-dimensional, and this turns out to be true for many other problem domains as well.

The following mechanism for the public project problem assumes that the designer can decide not only if the project will be built, but also which players will be allowed to use it. Thus, we now have many possible alternatives, that correspond to the different subsets of players that will be allowed to utilize the project. This is still a single-dimensional domain, as each player only cares about whether she is losing or winning, and so the alternatives, from the point of view of a specific player, can be divided to the two winning/losing

subsets. The following cost-sharing mechanism was proposed by Moulin [20] in a general cost-sharing framework. The mechanism is a direct-revelation mechanism, where each player, i , first submits her winning value, v_i . The mechanism then continues in rounds, where in the first round all players are present, and in each round one or more players are declared losers and retire. Suppose that in a certain round x players remain. If all remaining players have $v_i \geq C/x$ then they are declared winners, and each one pays C/x . Otherwise, all players with $v_i < C/x$ are declared losers, and “walk out”, and the process repeats. If no players remain then the project is rejected.

Clearly, the cost sharing mechanism always recovers the cost of the project, if it is indeed accepted. But is it truthful? One can analyze it directly, to show that indeed the dominant strategy of each player is to declare her true winning value. Perhaps a better way is to understand a characterization of truthfulness for the general abstract setting of a single-dimensional domain. For simplicity, we will assume that we require mechanisms to be “normalized”, i.e. that a losing player will pay exactly zero to the mechanism. Now, a mechanism is said to be “value-monotone” if a winner that increases her value will always remain a winner. More formally, for all $v_i \in V_i$ and $v_{-i} \in V_{-i}$, if i is a winner in the declaration (v_i, v_{-i}) then i is a winner in the declaration (v'_i, v_{-i}) , for all $v'_i \geq v_i$. Note that a value-monotone mechanism casts a “threshold value” function $v_i^*(v_{-i})$ such that, for every v_{-i} , player i wins when declaring $v_i > v_i^*(v_{-i})$, and loses when declaring $v_i < v_i^*(v_{-i})$. Quite interestingly, this structure completely characterizes incentive compatibility in single-dimensional domains:

Theorem 7 *A normalized direct-revelation mechanism for a single-dimensional domain is truthful if and only if it is value monotone and the price of a winning player is $v_i^*(v_{-i})$.*

Proof The first observation is that the price of a winner cannot depend on her declaration, v_i (only on the fact that she wins, and on the declaration of the other players). Otherwise, if it can depend on her declaration, then there are two possible bids v_i and v'_i such that i wins with both bids and pays p_j and p'_i , where $p'_i < p_i$. But then if the true value of

i is v_i then bidding v'_i instead of v_i will increase i 's utility, contradicting truthfulness.

We now show that a truthful mechanism must be value-monotone. Assume by contradiction that a declaration of (v_i, v_{-i}) will cause i to win, but a declaration of (v'_i, v_{-i}) will cause i to lose, for some $v'_i > v_i$. Suppose that i pays p_i for winning (when the others declare v_{-i}). Since we assume a normalized mechanism, truthfulness implies that $p_i \leq v_i$. But then when the true type of a player is v'_i , her utility from declaring the truth will be zero (she loses), and she can increase her utility by declaring v_i , which will cause her to win and to pay p_i , a contradiction.

Thus, a truthful mechanism must be value-monotone, and there exists a threshold value $v_i^*(v_{-i})$. To see that this defies p_i , let us first check the case of $p_i < v_i^*(v_{-i})$. In this case, if the type of i is v_i with $p_i < v_i < v_i^*(v_{-i})$, she will lose (by the definition of $v_i^*(v_{-i})$), and by bidding some false large enough v'_i she can win and get a positive utility of $v_i - p_i$. On the other hand, if $p_i > v_i^*(v_{-i})$ then with type v_i such that $p_i > v_i > v_i^*(v_{-i})$ a player will have negative utility of $v_i - p_i$ from declaring the truth, and she can strictly increase it by losing, again a contradiction. Therefore, it must be that $p_i = v_i^*(v_{-i})$.

To conclude, it only remains to show that a value-monotone mechanism with a price for a winner $p_i = v_i^*(v_{-i})$ is indeed truthful. Suppose first that with the truthful declaration i wins. Then $v_i > v_i^*(v_{-i}) = p_i$ and i has a positive utility. If she changes the declaration and remains a winner, her price does not change, and if she becomes a loser her utility decreases to zero. Thus, a winner cannot increase her utility. Similarly, a loser can change her utility only by becoming a winner, i.e. by declaring $v'_i > v_i^*(v_{-i}) > v_i$, but since she will then pay $v_i^*(v_{-i})$ her utility will now decrease to be negative. Thus, a loser cannot increase her utility either, and the mechanism is therefore truthful. $\square$

This structure of truthful mechanisms is very powerful, and reduces the mechanism design problem to the algorithmic problem of designing monotone social choice functions. Another strong implication of this structure is the fact that the payments of a truthful mechanism are completely derived from the social choice rule. Consequently,

if two mechanisms always choose the same set of winners and losers, then the revenues that they raise must also be equal. Myerson [21] was perhaps the first to observe that, in the context of auctions, and named this the “revenue equivalence” theorem.

As a result of this characterization, one can easily verify that the above-mentioned cost-sharing mechanism is indeed truthful. It is not hard to check that the two conditions of the theorem hold, and therefore its truthfulness is concluded. This is just one example of the usefulness of the characterization.

Multi-dimensional Domains

In the more general case, when the domain is multi-dimensional, the simple characterization from above does not fit, but it turns out that there exists a nice generalization. We describe two properties, cyclic monotonicity (Rochet [25]) and weak monotonicity (Bikhchandani et al. [7]), which achieve that. The exposition here also relies on [14]. It will be convenient to use the abstract social choice setting described above: there is a finite set A of alternatives, and each player has a type (valuation function) $v : A \rightarrow \mathfrak{R}$ that assigns a real number to every possible alternative. $v_i(a)$ should be interpreted as i 's value for alternative a . The valuation function $v_i(\cdot)$ belongs to the domain V_i of all possible valuation functions.

Our goal is to implement in dominant strategies the social choice function: $V_1 \times \cdots \times V_n \rightarrow A$. As before, it is not hard to verify that the required price function of a player i may depend on her declaration *only* through the choice of the alternative, i.e. that it takes the form $p_i : V_{-i} \times A \rightarrow \mathfrak{R}$, for every player i . For truthfulness, these prices should *satisfy* the following property. Fix any $v_{-i} \in V_{-i}$, and any $v_i, v'_i \in V_i$. Suppose that $f(v_i, v_{-i}) = a$ and $f(v'_i, v_{-i}) = b$. Then it is the case that:

$$v_i(a) - p_i(a, v_{-i}) \geq v_i(b) - p_i(b, v_{-i}) \quad (1)$$

In other words, player i 's utility from declaring his true v_i is no less than his utility from declaring some lie, v'_i , *no matter what the other players declare*. Given a social choice function f , the underlying question is what conditions should it satisfy to guarantee the existence of such prices.

Fix a player i , and fix the declarations of the others to v_{-i} . Let us assume, without loss of generality, that f is onto A (or, alternatively, define A' to be the range of $f(\cdot, v_{-i})$, and replace A with A' for the discussion below). Since the prices of Eq. (1) now become constant, we simply seek an assignment to the variables $\{p_a\}_{a \in A}$ such that $v_i(a) - v_i(b) \geq p_a - p_b$ for every $a, b \in A$ and $v_i \in V_i$ with $f(v_i, v_{-i}) = a$. This motivates the following definition:

$$\delta_{a,b} \doteq \inf \{v_i(a) - v_i(b) \mid v_i \in V_i, f(v_i, v_{-i}) = a\} \quad (2)$$

With this we can rephrase the above assignment problem, as follows. We seek an assignment to the variables $\{p_a\}_{a \in A}$ that satisfies:

$$p_a - p_b \leq \delta_{a,b} \quad \forall a, b \in A \quad (3)$$

By adding the two inequalities $p_a - p_b \leq \delta_{a,b}$ and $p_b - p_a \leq \delta_{b,a}$ we get that a necessary condition to the existence of such prices is the inequality $\delta_{a,b} + \delta_{b,a} \geq 0$. Note that this inequality is completely determined by the social choice function. This condition is termed the non-negative 2-cycle requirement. Similarly, for any k distinct alternatives $a_1, \dots, a_k$ we have the inequalities

$$\begin{aligned} p_{a_1} - p_{a_2} &\leq \delta_{a_1, a_2} \\ &\vdots \\ p_{a_{k-1}} - p_{a_k} &\leq \delta_{a_{k-1}, a_k} \\ p_{a_k} - p_{a_1} &\leq \delta_{a_k, a_1} \end{aligned}$$

and we get that any k -cycle must be non-negative, i.e. that $\sum_{i=1}^k \delta_{a_i, a_{i+1}} \geq 0$, where $a_{k+1} \equiv a_1$. It turns out that this is also a sufficient condition:

Theorem 8 *There exists a feasible assignment to (3) if and only if there are no negative-length cycles.*

One constructive way to prove this is by looking at the allocation graph”: this is a directed weighted graph $G = (V, E)$ where $V = A$ and $E = A \times A$, and an edge $a \rightarrow b$ (for any $a, b \in A$) has weight $\delta_{a,b}$. A standard basic result

of graph theory states that there exists a feasible assignment to (3) if and only if the allocation graph has no negative-length cycles. Furthermore, if all cycles are non-negative, the feasible assignment is as follows: set p_a to the length of the shortest path from a to some arbitrary fixed node $a^* \in A$.

With the above theorem, we can easily state a condition for implementability:

Definition 4 (Cycle Monotonicity) Social choice function f satisfies cycle monotonicity if for every player i , $v_{-i} \in V_{-i}$, some integer $k \leq |A|$, and $v_i^1, \dots, v_i^k \in V_i$

$$\sum_{j=1}^k \left[v_i^j(a_j) - v_i^j(a_{j+1}) \right] \geq 0$$

where $a_j = f(v_i^j, v_{-i})$ for $1 \leq j \leq k$, and $a_{k+1} = a_1$.

Theorem 9 f satisfies cycle monotonicity if and only if there are no negative cycles.

Corollary 1 A social choice function f is dominant-strategy implementable if and only if it satisfies cycle monotonicity.

This interesting structure implies, as another corollary, the fact that the prices are uniquely determined by the social choice function, for every connected domain (this was discussed above for the special case of single-dimensional domains). Very briefly, from the above, it follows that any two alternatives with $\delta_{ab} + \delta_{ba} = 0$ have $p_a - p_b = \delta_{ab} = -\delta_{ba}$. Thus, determining the price of one alternative completely determines the price of the second alternative. A short argument that we omit shows that the connectedness of the domain implies that for any two alternatives a and b , there's a path $a_1, \dots, a_k$ (with $a_1 = a$ and $a_k = b$) such that $\delta_{a_i, a_{i+1}} + \delta_{a_{i+1}, a_i} = 0$ for every $1 \leq i < k$. Thus, fixing the price of one alternative completely determines the prices of all other alternatives. In particular, there exists one alternative whose price is normalized to be (always) zero, then all other prices have also been completely

determined by the δ_{ab} 's weights (who in turn are completely determined by the function f).

Cycle monotonicity satisfies our motivating goal: a condition on f that involves only the properties of f , without existential price qualifiers. However, it is quite complex. k could be large, and a "shorter" condition would have been nicer. "Weak monotonicity (W-MON)" is exactly that:

Definition 5 (Weak Monotonicity) A social choice function f satisfies W-MON if for every player i , every v_{-i} , and every $v_i, v_i' \in V_i$ with $f(v_i, v_{-i}) = a$ and $f(v_i', v_{-i}) = b$, $v_i'(b) - v_i(b) \geq v_i'(a) - v_i(a)$.

In other words, if the outcome changes from a to b when i changes her type from v_i to v_i' then i 's value for b has increased at least as i 's value for a in the transition v_i to v_i' . W-MON is equivalent to cycle monotonicity with $k = 2$, or, alternatively, to the requirement of no negative 2-cycles. Hence it is necessary for truthfulness. As it turns out, it is also a sufficient condition on many domains. Very recently, Monderer [19] shows that weak monotonicity must imply cycle monotonicity if and only if the closure of the domain of valuations is convex. Thus, for such domains, it is enough to look at the more simple condition of weak monotonicity.

The Implementability of Non-Welfare-Maximizing Social Goals Now that the conditions for implementability are completely understood, it should be asked what forms of social choice functions satisfy them. We already saw that the welfare-maximizer function satisfies them, for any domain, and we ask what *other* implementable functions exist? For the single-dimensional case, we saw another example of a truthful mechanism, and the literature contains many more. For the multi-dimensional case, "interesting" examples are more rare, and a beautiful result by Roberts [25] shows that when the domain has full dimensionality then only weighted welfare maximizers are implementable. In other words, weak monotonicity implies welfare maximization. More precisely, a function f is an "affine maximizer" there exist weights $k_1, \dots, k_n$ and $\{C_x\}_{x \in A}$ such that, for all $v \in V$,

$$f(v) \in \arg \max_{x \in A} \left\{ \sum_{i=1}^n k_i v_i(x) + C_x \right\}$$

Roberts [25] shows that, if $|A| \geq 3$ and $V_i = \mathcal{R}^A$ for all i , then f is dominant-strategy implementable if and only if it is an affine maximizer.

However, most interesting domains are restricted in some meaningful way, and for this wide intermediary range of domains the current knowledge is rather scarce. One impossibility result that extends the result of Roberts to a restricted multi-dimensional case is given by Lavi et al. [18], who study multi-item auctions. In a multi-item auction, one seller (the mechanism designer) wants to allocate items to players (i.e. an alternative is an allocation of the items to the players). Lavi et al. [18] shows that every social choice function for multi-item auctions, that additionally satisfy four other social choice properties, must be an affine maximizer.

Before concluding the discussion on dominant-strategy implementation, we demonstrate the necessity for non-welfare-maximizers by considering the following “scheduling domain”. A designer wishes to assign n tasks/jobs to m workers, where worker i needs t_{ij} time units to complete job task j , and incurs a cost of t_{ij} for its processing time (one dollar per time unit). Importantly, this cost is private information of the worker, and workers are assumed to be strategic, each one selfishly trying to minimize its own cost. The load of worker i is the sum of costs of the jobs assigned to her, and the maximal load over all workers (in a given schedule) is termed the “makespan” of the schedule. The welfare maximizing social goal would put each task on the most efficient worker (for that task), which may result in a very high makespan. For example, consider a setting with two workers and n tasks. The first worker incurs a cost of 1 for every task, and the second worker incurs a cost of $1 + \varepsilon$ for every task. The social welfare is the minus of the sum of the costs of the two workers, and the VCG mechanism will therefore assign all tasks to the first worker. This is a very highly unbalanced allocation, which takes twice the time that the workers optimally need in order to finish all tasks (roughly splitting the work among them).

Thus, one may wish to consider a social goal different from welfare maximization, namely

makespan minimization. This goal aims to construct a balanced allocation, in order to minimize the completion time of the last task. Such an allocation can also be viewed as being a more “fair” allocation, in the sense of Rawls’ maxmin fairness criteria. Because of the strategic nature of the workers, we wish to design a truthful mechanism. While VCG is truthful, its outcome may be far from optimal, as demonstrated above. Nisan and Ronen [23], who have first studied this problem in the context of mechanism design, observed that VCG provides only an m -approximation” to the optimal makespan, meaning that VCG may sometimes produce a makespan that is m times larger than the optimal makespan. More importantly, they have shown that *no truthful deterministic mechanism can obtain an approximation ratio better than 2*. To date, the question of closing this gap between m and 2 remains open.

Archer and Tardos [1], on the other hand, considered a natural restriction this domain, that makes it single-dimensional, and showed with this they can construct many possibilities (for example, a truthful *optimal* mechanism). Thus, here too we see the contrast between single-dimensionality and multi-dimensionality. Lavi and Swamy [17] suggest a multi-dimensional special case, and give a truthful 2-approximation for the special case where the processing time of each job is known to be either “low” or “high”. This special case keeps the multi-dimensionality of the domain. The construction of this result does not rely on explicit prices, but rather uses the cycle-monotonicity condition described above, to construct a monotone allocation rule.

Budget Balancedness and Bayesian Mechanism Design

The previous sections portray a concrete picture of the advantages and the disadvantages of the solution concept of truthfulness in dominant strategies. On the one hand, this is a strong and convincing concept, which admits many positive results. However, there are several problems to all these results, that cannot be solved by a truthful mechanism. Among these, the budget-imbalance

problem was briefly mentioned, and this section looks again at this problem, as a motivation to the definition of the Bayesian Nash solution concept.

To recall the budget-imbalance problem of the VCG mechanism, let us consider a specific input to the Clarke mechanism from section “[Quasi-Linear Utilities and the VCG Mechanism](#)”: suppose the cost of the project is \$100, and there are 102 players, each values the project by \$1. It is a simple exercise to check that the Clarke mechanism will indeed choose to perform the project, and that each player will pay a price of zero (since the project would have been conducted even if a single player is removed). Thus, the mechanism designer does not cover the project’s cost. As described above, this problem, for this specific domain, can be fixed by considering the cost-sharing mechanism discussed in section “[The Importance of the Domain’s Dimensionality](#)”. However, this mechanism may sometimes choose not to perform the project although the society as a whole will benefit from performing it (i.e. it is not “socially efficient”), and, even more importantly, it is a solution only for the concrete domain of a public project. Is there a general mechanism (in the sense that VCG is general) that is both socially efficient and budget-balanced? In this section we describe such a mechanism, that was independently discovered by d’Aspremont and Ge’rardVaret [10] and by Arrow [3]. Its incentive compatibility will not be in dominant strategies. Instead, it is assumed that player types are drawn i.i. d. from some fixed and known cumulative distribution function F (the assumption that the types are drawn from the same distribution is not important, and is made here only for the ease of notation; the assumption that types are not correlated is important and cannot be removed in general). The solution concept of a Bayesian-Nash equilibrium is a natural extension of the regular Nash equilibrium concept, for a setting in which the distribution F is known to all players (this is termed the “common-prior” assumption), and where players aim to maximize the expectation of their quasi-linear utility.

Definition 6 A direct mechanism $M = (f, p)$ is Bayesian incentive compatible if for every player i , and for every $v_i, v'_i \in V_i$,

$$\begin{aligned} E_{v_{-i}}[v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i})] \\ \geq E_{v_{-i}}[v_i(f(v'_i, v_{-i})) - p_i(v'_i, v_{-i})] \end{aligned}$$

In other words, Bayesian incentive compatibility requires that a player will maximize her expected utility by declaring her true type. An alternative formulation is that truth-fulness in a Bayesian incentive compatible mechanism should be a “Bayesian-Nash equilibrium” (where the formal equilibrium definition naturally follows the above definition). This is an “ex-interim” equilibrium: the type of the player is already known to her, and the averaging is over the types of the others. A weaker equilibrium notion would be an “ex-ante” notion, where the player should decide on a strategy before knowing her own type, and so the averaging is done over her own types as well. A stronger notion would be an “ex-post” notion, where no-averaging is done at all, and the above inequality is required for every realization of the types of the other players. It can be shown that this stronger ex-post condition is equivalent to the requirement of dominant-strategy incentive compatibility. As a Bayesian-Nash equilibrium only considers the average over all possible realizations, it is clearly a weaker requirement than dominant-strategy implementability.

We will demonstrate the usefulness of this weaker notion by describing a general mechanism that is both ex-post socially efficient and ex-post budget balanced, and is Bayesian incentive-compatible. Define,

$$x_i(v_i) = E_{v_{-i}} \left[\sum_{j \neq i} v_j(f(v_i, v_{-i})) \right]$$

The “budget-balanced” (BB) mechanism asks the players to report their types, and then chooses the welfare-maximizing allocation according to the reported types (as VCG does). It then charges some payment $p_i(v_i, v_{-i}) = -x_i(v_i) + h_i(v_{-i})$, for some function $h_i(\cdot)$ that will be chosen later on in a specific way that balances the budget. But let us first verify that the mechanism is Bayesian incentive compatible, regardless of the choice of the functions $h_i(\cdot)$. Note that, for any realization of v_{-i} , we have that,

$$v_i(f(v_i, v_{-i})) + \sum_{j \neq i} v_j(f(v_i, v_{-i})) \geq v_i(f(v'_i, v_{-i})) + \sum_{j \neq i} v_j(f(v'_i, v_{-i}))$$

as the mechanism chooses the maximal-welfare alternative for the given reports. Clearly, taking the expectation on both sides will maintain the inequality. Therefore we get:

$$\begin{aligned} E_{v_{-i}}[v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i})] \\ &= E_{v_{-i}}[v_i(f(v_i, v_{-i}))] + E_{v_{-i}} \left[\sum_{j \neq i} v_j(f(v_i, v_{-i})) \right] \\ &\quad + E_{v_{-i}}[h_i(v_{-i})] \\ &\geq E_{v_{-i}}[v_i(f(v_i, v_{-i}))] + E_{v_{-i}} \left[\sum_{j \neq i} v_j(f(v_i, v_{-i})) \right] \\ &\quad + E_{v_{-i}}[h_i(v_{-i})] \\ &= E_{v_{-i}}[v_i(f(v'_i, v_{-i})) - p_i(v'_i, v_{-i})] \end{aligned}$$

which proves Bayesian incentive compatibility. To balance the budget, consider the specific function, $h_i(v_{-i}) = 1/(n-1) \sum_{j \neq i} x_j(v_j)$. Notice that the term $x_j(v_j)$ appears $(n-1)$ times in the sum $\sum_{i=1}^n h_i(v_{-i})$ for any $j = 1, \dots, n$. Therefore $\sum_{i=1}^n h_i(v_{-i}) = 1/(n-1) \sum_{j=1}^n (n-1)x_j(v_j) = \sum_{i=1}^n x_i(v_i)$. To conclude, we have $\sum_{i=1}^n p_i(v_i, v_{-i}) = \sum_{i=1}^n h_i(v_{-i}) - \sum_{i=1}^n x_i(v_i) = 0$, and the budget balancedness follows.

It is worth noting that such an exercise cannot be employed for the VCG mechanism, as there the “parallel” $x_i(\cdot)$ term should depend on the entire vector of declarations, not only on i ’s own declarations. This is the exact point where the averaging of the others’ valuations is crucial.

In addition to the difference in the solution concept, one other important advantage of VCG, in comparison with the BB mechanism, is the fact that VCG (with the Clarke payments) is ex-post “individually rational if a player declares her true valuation, it is guaranteed that she will not pay more than her value, no matter what the others will declare. Here, on the contrary, there is no reason why this should be true, in general. Can the solution concept of Bayesian incentive compatibility be used to construct a general budget-balanced and individually rational mechanism? In an important and influencing result, Myerson and

Satterthwaite [22] have shown that this is impossible: there is no general mechanism that satisfies the four properties (1) Bayesian incentive compatibility, (2) budget balancedness, (3) individual rationality, and (4) social efficiency. The proof uses a simple, natural exchange setting, where two traders (one buyer and one seller) wish to exchange an item. The seller has a cost c of producing the item, and the buyer obtains a value v from receiving it. Myerson and Satterthwaite show that there is no Bayesian incentive compatible mechanism that decides to perform the exchange if and only if $v > c$, such that Bayesian incentive compatibility and individual rationality are maintained, and the price that the buyer pays exactly equals the payment that the seller gets. In particular, VCG violates this last property, while BB satisfies it, but violates individual rationality (i.e. for some realizations of the values, a buyer may pay more than her value, or the seller may get less than her cost).

Besides this disadvantage of the BB mechanism, there are also additional disadvantages that result from the underlying assumptions of the solution concept itself. In particular, Bayesian incentive compatibility entails two strong assumptions about the characteristics of the players. First, it assumes that players are risk-neutral, i.e. care only about maximizing the expectation of their profit (value minus price). Thus, when players dislike risk, for example, and prefer to decrease the variance of the outcome, even on the expense of lowering the achieved expected profit, the rational of the Bayesian-Nash equilibrium concept breaks down. Second, the assumption of a common-prior, i.e. that all players agree on the same underlying distribution, seems strong and somewhat unrealistic. Often, players have different estimations about the underlying statistical characteristics of the environment, and this concept does not handle this well. Note that the solution concept of dominant-strategies does not suffer from any of these problems, which strengthens even more its importance. Unfortunately, the classical economics literature mainly ignores these disadvantages and problems. A well-known exception is the critique known as Wilson’s critique [28], who

raises the above mentioned problems, and argues in favor of “detail-free” mechanisms. Recently, this critique gained more popularity, and detail-free solution concepts are re-examined. For some examples, see [5, 6, 11].

Interdependent Valuations

Up to now, this entry described private value” models, i.e. models where the valuation (or the preference relation) of a player does not depend on the types of the other players. There are many settings in which this assumption is unrealistic, and a more suitable assumption is that the valuation of a specific player is affected by the valuations of the other players. This last statement may entail two interpretations. The first is that the distribution over the valuations of a specific player is correlated with the distribution over the valuations of the other players, and, thus, knowing a player’s actual valuation gives partial knowledge about the valuations of the other players. This first interpretation is still termed a private *value* model (but with correlated values instead of independent values), since after the player becomes aware of the actual realization of her valuation, she completely and fully knows her values for the different outcomes.

In contrast, with interdependent valuations, the actual valuation of a player depends on the actual valuations of the other players. Thus, a player does not fully know her own valuation. She only partially knows it, and can determine her full valuation only if given the others’ valuations as well. A classic example is a setting where a seller sells an oil field. The oil, of-course, is not seen on the ground surface, and the only way to exactly determine how much oil is there (and, by this, determine the actual worth of the field) is to extract it. Before buying the field, though, the potential buyers are only allowed to make preliminary tests, and by this to determine an estimation of the value of the field, which is not completely accurate. If all the buyers that are interested in the field have the same technical capabilities, it seems reasonable to assume that the *true value* of the field is the average over all the estimations

obtained by the different oil companies. Intuitively, a player that participates in an auction mechanism that determines who will buy the field, and at what price, has to act somehow as if she knows the value of the field, although she doesn’t. Clearly, this creates different complications. Such a model is very natural in auction settings, and indeed the entry on auctions handles the subject of interdependent valuations more broadly. Since this issue is also very relevant to general mechanism design theory, we describe here one specific, rather general result for mechanisms with interdependent valuations, to exemplify the definitions and the techniques being employed.

In the formal model of interdependent valuations, player i receives a signal $s_i \in S_i$, which may be multi-dimensional. Her valuation for a specific alternative $a \in A$ is a function of the signals $s_1, \dots, s_n$, i.e. $v_i: A \times S_1 \times \dots \times S_n \rightarrow \Re$. The case where $v_i(a, s_1, \dots, s_n) = v_j(a, s_1, \dots, s_n)$ for all players i, j and all $a, s_1, \dots, s_n$ is termed the “common value” case, as the actual values of all players are identical, and only their signals are different (as in the oil field example). The other extreme is when i ’s valuation depends only on i ’s signal, i.e. $v_i(a, s_1, \dots, s_n) = v_i(a, s_i)$, which is a return to the private value case. The entire range in general is termed the case of interdependent valuations. All the results described in the previous sections fail when we move to interdependent valuations. For example, in the VCG mechanism, a player is required to report her valuation function, which is not fully known to her in the interdependent valuation case. It turns out that the straight-forward modification of reporting the players’ signals does not maintain the truthfulness property, and, in fact, some strong impossibilities exist (Jehiel et al. [15]). However, interdependent valuations may also enable possibilities, and the classic result of Cremer and McLean [9] will be described here to exemplify this. This result shows how to use the interdependencies in order to increase the revenue of the mechanism designer, so that the entire surplus of the players can be extracted. Cremer and McLean [9] study an auction setting where there is one item for sale. n bidders have interdependent

values for the item, and it is assumed that the signal that each player receives is single-dimensional, i.e. each player receives a single real number as her signal. The valuation functions are assumed to be known to the mechanism designer, so that the only private information of the players are their signals. It is also assumed that the valuation functions are monotonically non-decreasing in the signals. For simplicity, it is assumed here that the signal space is discretized to be $S_i = \{0, \Delta, 2\Delta, \dots\}$. The last (and crucial) assumption is that the valuation functions satisfy the “single-crossing” property: if $v_i(s_i, s_{-i}) \geq v_j(s_i, s_{-i})$ then $v_i(s_i + \Delta, s_{-i}) \geq v_j(s_i + \Delta, s_{-i})$. This says that i 's signal affects i 's own value (weakly) more than it affects the value of any other player. This last assumption is strong, but in some sense necessary, as it is possible to construct interdependent valuation functions (that violate single-crossing) for which no truthful mechanism can be efficient (i.e. allocate the item to the player with the highest value).

Consider the following CM mechanism for this problem: each player reports her signal, and the player with the highest *value* (note that this may be different than the player with the highest signal) receives the object. In order to determine her payment, define the “threshold signal” $T_i(s_{-i})$ of any player i to be the minimal signal that will enable her to win (given the signals of the other players), i.e. $T_i(s_{-i}) = \min \{\tilde{s}_i \in S_i \mid v_i(\tilde{s}_i, s_{-i}) \geq \max_{j \neq i} v_j(\tilde{s}_i, s_{-i})\}$. The payment of the winner, i , is her value if her signal was $T_i(s_{-i})$, i.e. $P_i(s_{-i}) = v_i(T_i(s_{-i}), s_{-i})$. Clearly, if all players report their true signals, then the player with the highest value receives the item. Truthful reporting is also an ex-post Nash equilibrium, which means the following: if all other players report the true signal (no matter what that is) then it is a best response for i to report her true signal as well.

To verify that truthfulness is indeed an ex-post Nash equilibrium, notice first that each player has a price for winning which does not depend on her declaration. Now, truthful reporting will ensure winning (given that the others are truthful as well) if and only if the true value of the player is

higher than her price (i.e. iff winning will yield a positive utility). Thus, when a player “wants to win”, truthful reporting will do that, and when a player “wants to lose”, truthful reporting will do that as well, and so truthfulness will always maximize the player's utility.

The notion of an ex-post equilibrium is stronger than Bayesian-Nash equilibrium, since, here, even after the signals are revealed no player regrets her declaration (while in Bayesian-Nash equilibrium, since only the *expected* utility is maximized, there are some realizations for which a player can deviate and gain). On the other hand, ex-post equilibrium is weaker than dominant strategies, in which truthfulness is the best strategy no matter what the others choose to declare, while here truthfulness is a best response only if the others are truthful as well.

As seen above, both for the VCG mechanism as well as for the BB mechanism, adding a “constant” to the prices (i.e. setting $\tilde{P}_i(s_{-i}) = P_i + h_i(s_{-i})$) maintains the strategic properties of the mechanism, since the function $h_i(\cdot)$ does not depend on the declaration of player i . The correlation in the values can help the mechanism designer to extract more payments from the players, as follows. Consider the matrix that describes the conditional probability for a specific tuple of signals of the other players, given i 's own signal. There is a row for every signal s_i of i , a column for every tuple of signals s_{-i} of the other players, and the cell (s_i, s_{-i}) contains the conditional probability $Pr(s_{-i} | s_i)$. In the private value case, the signals of the players are not correlated, hence the matrix has rank one (all rows are identical). As the correlation between the signals “increases”, the rank increases, and we consider here the case when the matrix has full row rank. Let $q_i(s_i, s_{-i})$ be an indicator to the event that i is the winner when the signals are (s_i, s_{-i}) . The expected surplus of player i in the CM mechanism is $U_i^*(s_i) = \sum_{s_{-i}} Pr(s_{-i} | s_i) \cdot (q_i(s_i, s_{-i}) \cdot v_i(s_i, s_{-i})) - P_i(s_{-i})$. ($P_i(s_{-i})$ is defined to be zero whenever i is not a winner). Now find “constants” $h_i(s_{-i})$ such that, for every s_i , $\sum_{s_{-i}} h_i(s_{-i}) \cdot Pr(s_{-i} | s_i) = U_i^*(s_i)$. Note that such an $h_i(\cdot)$ function exists: we have a system of linear equations, where the variables are the function values $h_i(s_{-i})$ for all possible tuples s_{-i} and the qualifiers are the probabilities and the expected

surplus. Since the matrix of qualities has full row rank, a solution exists. It is now not hard to verify that, with prices $\hat{P}_i(\cdot)$, the expected utility of a truthful player is zero.

As mentioned above, truthfulness is still an ex-post equilibrium of this mechanism. It is not ex-post individually rational, though, but rather only ex-ante, since a player pays her expected surplus even if the actual signals cause her to lose. Thus, this mechanism can be considered a fair lottery. Also note that the crucial property was the correlation between the values, the interdependence assumption was not important.

Future Directions

As surveyed here, the last three decades have seen the theory of mechanism design being developed in many different directions. The common thread of all settings is the requirement to implement some social goal in the presence of incomplete information—the social designer does not know the players' preferences for the different outcomes. We have seen several alternative assumptions about the structure of players' preferences, the different equilibria solution concepts that are suitable for the different cases, and several positive examples for elegant solutions. We have also discussed some impossibilities, demonstrating that some attractive definitions may turn out almost powerless. One relatively new research direction in mechanism design is the analysis of new models for the emerging Internet economy, and the development of new alternative solution concepts that better suit this setting. A very recent example is the new model of “dynamic mechanism design”, where the parameters of the problem (e.g. the number players, or their types) vary over time. Such settings become more and more important as the economic environment becomes more dynamic, for example due to the growing importance of the electronic markets. Examples for such models include e.g. the works by Lavi and Nisan [16] in the context of computer science models, and by Athey and Segal [4] in a more classical economic context, among many other works that study such dynamic settings.

The Internet environment also strengthens the question marks posed on the solution concept of Bayesian incentive compatibility, which was the most common solution concept in mechanism design literature in the 1980s and throughout the 1990s, due to the accompanying assumption of a common prior. Such an assumption seems problematic in general, and in particular in an environment like the Internet, that brings together players from many different parts of the world. It seems that the research community agrees more and more that alternative, detail-free solution concepts should be sought. The description of more recently new solution concepts is beyond the scope of this entry, and the interested reader is referred, for example, to the papers by [5, 6, 11] for some recent examples.

Another aspect of mechanism design that is largely ignored in the classic research is the computational feasibility of the mechanisms being suggested. This question is not just a technicality—some classic mechanisms imply heavy computational and communicational requirements that scale exponentially as the number of players increase, making them completely infeasible for even moderate numbers players. The computer science community has begun looking at the design of computationally efficient mechanisms, and the recent book by Nisan et al. [24] contains several surveys on the subject.

Bibliography

1. Archer A, Tardos E (2001) Truthful mechanisms for one-parameter agents. In: Proceedings of the 42nd Annual Symposium on Foundations of Computer Science, FOCS'01, Las Vegas. IEEE Computer Society
2. Arrow K (1951) Social Choice and Individual Values. Wiley, New York
3. Arrow K (1979) The property rights doctrine and demand revelation under incomplete information. In: Boskin M (ed) Economics and Human Welfare. Academic Press, New York
4. Athey S, Segal I (2007) Designing dynamic mechanisms. *Am Econ Rev* 9(2):131–136
5. Babaioff M, Lavi R, Pavlov E (2006) Single-value combinatorial auctions and implementation in undominated strategies. In: Proceedings of the 17th symposium on discrete algorithms, SODA, Miami. ACM Press
6. Bergemann D, Morris S (2005) Robust mechanism design. *Econometrica* 73:1771–1813

7. Bikhchandani S, Chatterjee S, Lavi R, Mu'alem A, Nisan N, Sen A (2006) Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica* 74(4): 1109–1132
8. Clarke E (1971) Multipart pricing of public goods. *Public Choice* 8: 17–33
9. Cremer J, McLean R (1985) Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica* 53:345–361
10. d'Aspremont C, Gérard-Varet L (1979) Incentives and incomplete information. *J Public Econ* 11:25–45
11. Dekel E, Wolinsky A (2003) Rationalizable outcomes of large private-value first-price discrete auctions. *Games Econ Behav* 43 (2): 175–188
12. Gibbard A (1973) Manipulation of voting schemes: A general result. *Econometrica* 41 (4):587–601
13. Groves T (1973) Incentives in teams. *Econometrica* 41 (4): 617–631
14. Gui H, Muller R, Vohra RV (2004) Characterizing dominant strategy mechanisms with multidimensional types. Working paper, unpublished
15. Jehiel P, Meyer ter Vehn M, Moldovanu B, Zame WR (2006) The limits of ex-post implementation. *Econometrica* 74(3):585–610
16. Lavi R, Nisan N (2004) Competitive analysis of incentive compatible on-line auctions. *Theor Comput Sci* 310:159–180
17. Lavi R, Swamy C (2007) Truthful mechanism design for multidimensional scheduling. In: The Proceedings of the 8th ACM conference on electronic commerce, EC'07, San Diego. ACM Press
18. Lavi R, Mu'alem A, Nisan N (2003) Towards a characterization of truthful combinatorial auctions. In: Proceedings of the 44rd annual symposium on foundations of computer science, FOCS'03, Cambridge. IEEE Computer Society
19. Monderer D (2007) Monotonicity and implementability. Working paper, unpublished
20. Moulin H (1999) Incremental cost sharing: Characterization by coalition strategy-proofness. *Soc Choice Welf* 16:279–320
21. Myerson R (1981) Optimal auction design. *Math Oper Res* 6: 58–73
22. Myerson R, Satterthwaite M (1983) Efficient mechanisms for bilateral trading. *J Econ Theor* 29:265–281
23. Nisan N, Ronen A (2001) Algorithmic mechanism design. *Games Econ Behav* 35:166–196
24. Nisan N, Roughgarden T, Tardos E, Vazirani W (eds) (2007) Algorithmic game theory. Cambridge University Press, New York
25. Roberts K (1979) The characterization of implementable choice rules. In: Laffont JJ (ed) Aggregation and revelation of preferences. North-Holland, Amsterdam, pp 321–349
26. Rochet JC (1987) A necessary and sufficient condition for rationalizability in a quasilinear context. *J Math Econ* 16:191–200. 27
27. Satterthwaite M (1975) Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *J Econ Theor* 10:187–217. 28
28. Vickrey W (1961) Counterspeculations, auctions, and competitive sealed tenders. *J Financ* 16:8–37
29. Wilson R (1987) Game-theoretic analyses of trading processes. In: Bewley T (ed) Advances in economic theory: fifth world congress. Cambridge University Press, New York, pp 33–70

Auctions

Martin Pesendorfer

Department of Economics, London School of
Economics and Political Science, London, UK

Article Outline

Glossary

Introduction

Second-Price Auction

First-Price Sealed-Bid Auction

Comparing Auction Outcomes

Empirics of Auctions

Winner's Curse

Collusive Bidding

Concluding Remarks

Bibliography

Glossary

All-pay first-price auction Bidders submit sealed bids. The high bidder wins the item. In case of a tie, the auctioneer randomly assigns the item to one of the high bidders with equal probability. All bidders (including losing bidders) pay their bid.

Bayesian Nash equilibrium A Bayesian Nash equilibrium is a collection of bidding strategies so that (i) no bidder has an incentive to deviate and (ii) beliefs are consistent with the underlying informational assumptions.

Bidding strategy A bidding strategy for a buyer is a mapping from the buyer's signal into bid prices.

Dutch auction Price falls until one bidder presses her button. That bidder gets the object at the current price. Losers pay nothing.

English auction Bidders call out successively higher prices until one bidder remains. The item is allocated to the last remaining bidder at the price at which the second last bidder dropped out.

First-price auction Bidders submit sealed bids. The high bidder wins the item and pays her bid. In case of a tie, the auctioneer randomly assigns the item to one of the high bidders with equal probability. Losers pay nothing.

Second-price auction Bidders submit sealed bids. The high bidder wins the item and pays second highest bid. In case of a tie, the auctioneer randomly assigns the item to one of the high bidders with equal probability. Losers pay nothing.

Introduction

Auctions have been a common selling form throughout history; see Cassidy (1967) for an account. Roman legions sold their plunder at auction. Slave auctions were held throughout medieval times. Art auctions have been taking place for the last 300 years with Christie's and Sotheby's being two well-known auction houses. Real estate, treasury bills, flowers, livestock, even large corporations are sold at auction. Government procurement follows specific regulations and rules which give rise to an auction rule. The sale of mineral extraction rights and spectrum licenses are good sources for governmental revenues. eBay has become a successful marketplace with the arrival of the Internet.

This entry surveys contributions of the auction literature. It is a selected account from an economist's perspective; see McAfee and McMillan (1987), Klemperer (1999), Krishna (2002), Hong and Paarsch (2006), and Hortacsu and McAdams (2016) for related surveys. Auctions are encountered in many settings, but specific rules and procedures may differ. Broadly speaking, we can distinguish single-item versus multi-item, sealed-bid versus open-outcry, and single-round versus multi-round auctions. The nature of the rules and format may depend on the items at hand but will also affect the behavior of bidders and what revenues the seller may get. Popular single-item auction rules include:

- **English** open-outcry auction in which bidders call out successively higher prices until one bidder remains. The item is allocated to the last remaining bidder at the price at which the second last bidder dropped out. Sometimes these are referred to as “hammer auctions” which are commonly used by Sotheby’s and Christy’s.
- **Second-price** sealed-bid auction in which bidders submit sealed bids. The high bidder wins the item and pays the second highest bid. In case of a tie, the auctioneer randomly assigns the item to one of the high bidders with equal probability. Losers pay nothing.
- **Dutch** or descending-price auction is the opposite of English auction. Price falls until one bidder presses her button. That bidder gets the object at the current price. Losers pay nothing. This auction format is used to sell flowers in Holland.
- **First-price** sealed-bid auction has bidders submitting sealed bids. The high bidder wins the item and pays her bid. In case of a tie, the auctioneer randomly assigns the item to one of the high bidders with equal probability. Losers pay nothing. The first-price auction format is commonly used for governmental procurement.
- **All-pay** first-price sealed-bid auction has bidders submitting sealed bids. The high bidder wins the item. In case of a tie, the auctioneer randomly assigns the item to one of the high bidders with equal probability. All bidders (including losing bidders) pay their bid.

The seller of the item selects the auction format before bidding starts. The seller may additionally decide how much information to reveal about the item and may announce a reserve price, which is a minimum price at which the seller is willing to sell. If the seller knows what buyers are willing to pay, then the seller can post a selling price equal to the high willingness to pay. The resulting allocation extracts all the rent and gives the item to the buyer that values it most. When the seller does not know the willingness to pay, then this posting price scheme may perform poorly. An auction seems then a good choice as the seller may achieve higher revenues when using an auction

than any posting price scheme, as was shown by Myerson (1981). An auction enables the seller to learn about buyers’ willingness to pay and to extract as much rent as possible.

The informational environment, that is, how much buyers know about their own willingness to pay, can vary. Environments range from buyers knowing precisely their value of the object or having a rough idea only. Let’s describe some informational structures. Suppose there are N buyers. Let $\theta_i \in [0, A]$ denote buyer i ’s signal drawn identically and independently from a cumulative distribution function F with probability density function f . Formally, the joint probability density function of signals equals the product of the marginals, $f(\theta_1, \dots, \theta_N) = \prod_i f(\theta_i)$ and this is a

common knowledge. The independence assumption is for illustration purposes only. It can be relaxed leading to a correlated information environment; see Milgrom and Weber (1982) for affiliation which is a particular correlation structure among signals. Heterogeneity between buyers can be incorporated by allowing signal distributions to differ across bidders, $F_1, \dots, F_N$. Following standard incomplete information terminology, we assess information at the interim stage in which buyer i has learned her own signal θ_i but has not learned the value of competitors’ signals $\theta_j, j \neq i$ (In information economics, the term *interim* is used to distinguish from *ex ante*, in which types are not known yet, and *ex post*, in which everyone knows everyone else’s type.).

Private values arise when the signal equals the value of the object to buyer i , $v_i = \theta_i$. Private values refer to a situation in which buyers know precisely their own value. An example may be a construction contract in which firms know their own opportunity costs of undertaking the project. However, a firm may be only vaguely informed about competing firms’ costs.

(Pure) common values arise when the value of the object is determined by the average signal, $v = \frac{1}{N} \sum_j \theta_j$. This environment differs from private values in two respects: first, as buyer i only observes one signal θ_i from a total of N signals, she knows only a little bit about the true value and,

second, all bidders assign the same value v . An example may be an oil field auction. Each bidder conducts their own study of how much oil there is and comes up with an estimate θ_i . The true value of the oil field will be some average across the bidders' noisy signals and is the same to all bidders. A second example is the wallet game, in which two bidders compete for the joint value of their wallets. In that case the total value equals the sum of signals, $v = \theta_1 + \theta_2$, or twice the average signal.

Interdependent values are a mixture between private and common values. With interdependent values, the value of the object to buyer i is given by $v_i = \alpha\theta_i + (1 - \alpha)\frac{1}{N-1}\sum_{j \neq i}\theta_j$ with $0 < \alpha < 1$. When the parameter $\alpha = 1$, this formula reduces to $v_i = \theta_i$, the case of private values. On the other hand, when the parameter $\alpha = \frac{1}{N}$, then this formula becomes the pure common value case.

Most of our following analysis will focus on the case of (independent) private values.

Assumption Buyers know their values privately, $v_i = \theta_i$ for all i .

The reader interested in the more advanced topic of interdependent valuations is referred to Krishna (2002) for a nice introductory exposition. We shall illustrate on occasions what may happen under alternative informational assumptions.

The seller may have information that is useful to bidders. For example, on eBay, the seller may decide how accurately to describe the object. Or a used car owner may know very well the pros and cons of the car. What should the seller do? Reveal the information prior to bidding, or conceal it. The seller may also impose a reserve price. Should the seller impose a reserve price? At what level? We shall return to these questions in the section entitled "Comparing Auction Outcomes" after having studied buyer's behavior in standard auction formats.

The sections entitled "Second-Price Auction" and "English Auction" examine optimal bidding strategies and Bayesian Nash equilibria for standard auction formats. The payoff bidder i receives will depend on the auction rule, the equilibrium played, and attitudes toward risk. Attitudes toward risk matter as bidders face lotteries, winning the auction or not. We shall assume risk neutrality in which bidder i 's payoff equals the

expected value of the lottery. We shall comment later on the extension to bidder risk aversion is introduced.

The strategy space and equilibrium concept is shared across the following sections.

A **bidding strategy** for buyer i is a mapping from signals into bid prices, $b_i : [0, A] \rightarrow \mathbb{R}$.

A Bayesian Nash equilibrium is a collection of bidding strategies, $(b_i)_{i=1}^N$, so that (i) no bidder has an incentive to deviate and (ii) beliefs are consistent with the underlying informational assumptions.

With a Bayesian Nash equilibrium, we mean a stable resting point in which every bidder adopts a strategy that maximizes her payoff.

The entry is organized as follows: We shall begin by studying bidder behavior at specific auction rules, including second-price and first-price auction. We then compare bidders' payoffs and auctioneer's revenues across distinct auction formats. We describe key issues for empirical work on auctions, the winner's curse phenomenon, and issues concerning collusive behavior at auction.

Second-Price Auction

The second-price auction format has been advocated in Vickrey (1961). It has the rule that the high bidder wins the item and pays the second highest bid. The payoff for a winning buyer i is $v_i - b_{(2)}$ where $b_{(2)}$ is the second highest bid. The payoff to a losing bidder is zero.

What is an optimal bidding strategy in a second-price auction? Let's consider an example. Suppose your value is 60; what should you bid? You could bid your value. Is it optimal to bid your value? The answer is yes. Suppose another bidder bids 70. Do you regret? No as you would lose money if you outbid the other bidder. Suppose other bidders' bid is 40. Do you regret? No. So, bidding your value is indeed optimal. It is a Nash equilibrium.

This example generalizes and leads us to the following result.

Theorem 1 Bidding the true value, $b_i(v_i) = v_i$, is a Bayesian Nash equilibrium in weakly dominant strategies.

Proof Let the seller's reserve price be denoted by R . Suppose buyer i 's valuation is below the reserve price, $v_i < R$: It is (weakly) optimal for buyer i to bid v_i . The only way that buyer i can win the item is if buyer i bids more than i , but in case of winning, she pays at least R and makes a loss, $v_i - R < 0$. Next, suppose buyer i 's valuation is above the reserve price, $v_i \geq R$: Following the strategy implies that if buyer i wins, she pays the second highest bid $b_{(2)} < v_i$ and makes a profit of $v_i - b_{(2)}$. Consider a deviation: (a) Suppose buyer i bids more than her valuation, $b_i > v_i$. For $b_{(2)} \leq v_i$, she pays $b_{(2)}$ and she gets the same payoff, but for $b_i > b_{(2)} > v_i$, she makes a loss. (b) Suppose buyer i bids less than her valuation, $b_i \leq v_i$. For $b_{(2)} < b_i$, she wins, pays $b_{(2)}$, and gets the same payoff, but for $b_i \leq b_{(2)} < v_i$, she does not win the auction, gets a payoff that is zero, and makes a loss.

The theorem characterizes an equilibrium which has the feature that all bidders bid their true value, bidders bid "sincerely." An interesting feature of the second-price auction is that sincere bidding is optimal irrespective what bidding strategies other bidders adopt.

Vickrey (1961) is the classic auction paper that has emphasized these features of second-price auctions and compares the second-price auction to a first-price sealed-bid auction. Vickrey has shown that the above equilibrium is in fact efficient. With efficiency we mean that the bidder that values the item the most gets it.

We can also assess the revenues to the seller. The expected revenues (for the seller) of the Vickrey auction equal the expected second highest valuation.

Are there other equilibria in second-price auctions? The answer is yes. Suppose there is no reserve price. Consider the following "pooling equilibrium" in which bidder 1 bids $b_1 = A$ and all other bidder bid $b_i = 0$. This is in fact a Bayesian Nash equilibrium. Nobody can benefit from deviating. Notice though that this is not a dominant strategy equilibrium. Moreover, the outcome is not efficient. We shall focus on the dominant strategy equilibrium when comparing auction outcomes across auction formats and return to the pooling equilibrium in the section describing empirical work.

Next, we shall illustrate that the equilibrium in the English auction shares features with the above equilibria.

English Auction

We consider a continuous-price version of the English auction in which the price increases continuously, without any bidding jumps, and does so until only one bidder remains. As the price increases, bidders drop out irrevocably. When only one bidder remains, the item is allocated to the last remaining bidder at the price at which the second last bidder dropped out.

Consider a bidding strategy in the English auction in which bidder i stays in until the price reaches her value v_i . If everyone adopts this strategy, does this constitute an equilibrium? Yes, for the same reason as given in the above proof.

In fact there is a strategic equivalence between an English auction and a second-price auction with independent private information. Think of an agent that bids on behalf of bidder i . The agent would receive a number to submit in a second-price auction and a dropout value in the English auction. With strategic equivalence, it is meant that the number should be identical in both auction formats.

Notice though that the strategic equivalence breaks down with interdependent values. In that case, we need to take into account that bidders form their valuation estimate based on all available information. In a second-price sealed-bid auction, the only available information is a bidder's private signal. In an English auction, bidders learn something as the price increases. For instance, when an opponent drops out of the auction, something can be inferred about that bidder's private signal which may influence the valuation estimate. Thus, as price increases, bidders will update their bidding strategy to take the additional information into account.

Next, we shall consider bidding behavior at a first-price auction.

First-Price Sealed-Bid Auction

In a **first-price** auction, bidders submit sealed bids. The high bidder wins the item and pays her

bid. In case of a tie, the auctioneer randomly assigns the item to one of the high bidders with equal probability. Losers pay nothing.

We maintain our assumption of bidder risk neutrality. If bidder i wins the item, her payoff is $v_i - b_i$, while she makes a payoff of zero if she loses. Ignoring the issue of ties for the moment, let $\Pr(p_i > b_j, \text{ for all } j)$ denote the winning probability. It is the probability that bidder i submits the high bid.

Bidder i 's (interim) expected payoff U_i will depend on the valuation v_i and the bid b_i submitted. For a first-price auction, the (interim) expected payoff equals

$$U_i(v_i, b_i) = [v_i - b_i] \cdot \Pr(b_i > b_j, \text{ for all } j) \quad (1)$$

Recall that a strategy is a mapping: $b_i : [0, A] \rightarrow \mathbb{R}$. We wish to find an equilibrium. One way to proceed is to invoke calculus and work through the first-order conditions. Another way to proceed is to impose assumptions on the equilibrium and use those in the derivation. At the end, we then have to verify that indeed it is an equilibrium. This is the approach we shall adopt.

We restrict attention to bidding strategies which are differentiable, strict monotone increasing, and symmetric. Thus, there exists a differentiable strict monotone increasing function $b : [0, A] \rightarrow \mathbb{R}$ so that

$$b_i(v_i) = b(v_i) \text{ for all } i \text{ and } v_i \in [0, A] \quad (2)$$

The restriction allows us to simplify the problem. The winning probability when other bidders use the strategy $b(v_j)$ is given by $\Pr(b_i > b_j, \text{ for all } j \neq i) = F(b^{-1}(b_i))^{N-1}$. To see this, consider the following equivalent expressions:

$$\begin{aligned} & \Pr(b_i > b_j, \text{ for all } j \neq i) \\ &= \Pr(b_i > b(v_1), \dots, b_i > b(v_N)) \\ &= \Pr(b^{-1}(b_i) > v_1) \cdots \Pr(b^{-1}(b_i) > v_N) \\ &= F(b^{-1}(b_i))^{N-1} \end{aligned}$$

The first line writes out what bidder i 's winning probability is when other bidders follow the bidding strategy $b(\cdot)$. The second equation uses the independence assumption and the strict monotonicity

property, which implies that the inverse function b^{-1} exists and is strict monotone. The final equation uses the definition of the cumulative distribution function F , which evaluated at the number $b^{-1}(b_i)$ gives the probability that a buyer's valuation is less than equal to that number. In total this probability arises $(N-1)$ times as there are $(N-1)$ opponents.

The (interim) expected payoff becomes thus

$$U_i(v_i, b_i) = [v_i - b_i] \cdot F(b^{-1}(b_i))^{N-1} \quad (3)$$

In a Bayesian Nash equilibrium, it must be that the strategy $b(v_i)$ is optimal. Observe that $b(0) = 0$ since the valuation is zero. Observe also that any bid greater than $b(A)$ will win for sure. Thus, possible deviation bids must be contained in the range $[0, b(A)]$. Since the bid strategy $b(\cdot)$ is strict monotone, we can express this range with $b(w)$ and $w \in [0, A]$.

If bidder i with valuation $x = v_i$ bids $b(w)$ rather than $b(x)$, her payoff is

$$U_i(x, b(w)) = [x - b(w)]F(w)^{N-1}$$

Taking the derivative with respect to w yields

$$\frac{\partial U_i(x, b(w))}{\partial w} = [x - b(w)] \frac{\partial (F(w)^{N-1})}{\partial w} - b'(w)F(w)^{N-1}$$

For bidding $b(x)$ to be optimal, this derivative must be zero when evaluated at $w = x$

$$x \left[\frac{\partial (F(x)^{N-1})}{\partial x} \right] = b(x) \left[\frac{\partial (F(x)^{N-1})}{\partial x} \right] + b'(x)F(x)^{N-1}$$

Integrating both sides with respect to x

$$\int_0^{v_i} x \left[\frac{\partial (F(x)^{N-1})}{\partial x} \right] dx = b(v_i)F(v_i)^{N-1} - b(0)F(0)^{N-1}$$

Since $b(0) = 0$ and $F(0) = 0$, we have an explicit solution for the bid function

$$b(v_i) = \frac{\int_0^{v_i} x \left[\frac{\partial (F(x)^{N-1})}{\partial x} \right] dx}{F(v_i)^{N-1}}$$

Observe that this strategy is indeed differentiable and strict monotone. It is thus an equilibrium. The right-hand side is the conditional expected value of the random variable $v_{(2)}$ with probability density function $\frac{\partial(F(x)^{N-1})}{\partial x}$ conditional on $v_{(2)}$ being less than the valuation v_i . This leads us to the following result for first-price auctions.

Theorem 2 (First-Price Auction Equilibrium)

The Bayesian Nash equilibrium bid function in the first-price auction is

$$b(v_i) = E[v_{(2)} | v_{(2)} < v_i]$$

The equilibrium bid function has the feature that bidder i marks her valuation v_i down. Bidder i bids the expected value of the high-valuation competitor conditional on her competitors' valuations being less than her own.

Example Let's consider a simple two bidder case with F the uniform distribution on $[0,1]$. The equilibrium bid function becomes

$$\begin{aligned} b(v_i) &= \frac{\int_0^{v_i} x dx}{v_i} \\ &= \frac{v_i}{2} \end{aligned}$$

What is $\frac{v_i}{2}$? It equals the expected valuation of your competitor conditional on your competitor's valuation being less than your own, $E[v_{(2)} | v_{(2)} < v_i]$.

Observe that the first-price auction is efficient. The bidder with the high valuation wins the item. The reason is that the equilibrium bid function is identical and strict monotone increasing.

The equilibrium derivation was based on symmetric, strict monotone bid functions. The questions arise whether there are other equilibria. The answer is no. Maskin and Riley (2003) have shown that the above bidding strategy is the only equilibrium in first-price auctions.

Richer informational environments have been studied by a number of authors. Milgrom and Weber (1982) is the classic reference for equilibria in standard auction formats with symmetric bidders with interdependent values and affiliated signals encompassing both the common value and private value models. Bergemann et al. (2017) examine bidding implications in first-price auctions when the information structures specifying bidders' information about their own and others' values are not restricted.

We next illustrate the Bayesian Nash equilibrium in two variants of the first-price auction: (i) in which all bidders pay their bid and (ii) in the Dutch auction.

All-Pay First-Price Auction

In an **all-pay first-price** sealed-bid auction bidders submit sealed bids. The high bidder wins the item. In case of a tie, the auctioneer randomly assigns the item to one of the high bidders with equal probability. All bidders (including losing bidders) pay their bid.

Bidder i 's (interim) expected payoff U_i will depend on her valuation v_i and her submitted bid v_i . For an all-pay first-price auction, the (interim) expected payoff equals

$$U_i(v_i, b_i) = v_i \cdot \Pr(b_i > b_j, \text{ for all } j) - b_i \quad (4)$$

Recall that a strategy is a mapping $b_i: [0, A] \rightarrow \mathbb{R}$. We wish to find an equilibrium. We follow a similar approach as in the first-price auction which leads to a differential equation. Suppose bidders use a strictly increasing bid function $b(\cdot)$. As before, this implies that $\Pr(b_i > b_j, \text{ for all } j) = F(b^{-1}(b_i))^{N-1}$. If bidder i with valuation $x = v_i$ uses $b(w)$ rather than $b(x)$, her payoff is

$$U_i(x, b(w)) = xF(w)^{N-1} - b(w)$$

Taking the derivative with respect to w yields

$$\frac{\partial U_i(x, b(w))}{\partial w} = x \frac{\partial (F(w)^{N-1})}{\partial w} - b'(w)$$

For bidding $b(x)$ to be optimal, this derivative must be zero when evaluated at $w = x$

$$x \left[\frac{\partial (F(w)^{N-1})}{\partial w} \right]_{w=x} = b'(x)$$

Integrating both sides with respect to x

$$\int_0^{v_i} x \left[\frac{\partial (F(w)^{N-1})}{\partial w} \right]_{w=x} dx = \int_0^{v_i} b'(x) dx$$

The left-hand side can be integrated by parts, which yields

$$\left[xF(x)^{N-1} \right]_0^{v_i} - \int_0^{v_i} F(x)^{N-1} dx = b(v_i) - b(0)$$

Since $b(0) = 0$, we have an explicit solution for the bid function

$$b(v_i) = v_i F(v_i)^{N-1} - \int_0^{v_i} F(x)^{N-1} dx$$

This leads us to the following result for all-pay first-price auctions.

Theorem 3 (All-Pay First-Price Auction Equilibrium) The Bayesian Nash equilibrium bid function in the all-pay first-price auction is $b(v_i) = v_i F(v_i)^{N-1} - \int_0^{v_i} F(x)^{N-1} dx$.

Observe that the all-pay first-price auction is efficient. The bidder with the high valuation wins the item. The reason is that the equilibrium bid function is identical and strict monotone increasing.

Dutch Auction

The **Dutch** auction has price falling until one bidder jumps in. That bidder gets the object at the current price. Losers pay nothing.

A Dutch auction is strategically equivalent to the first-price auction. Think of an agent bidding on behalf of a bidder. In both formats, the bidder would instruct the agent with a number to bid. The number is the same in both formats.

We have now considered a number of auction rules and formats. Next, we shall compare the outcomes under those auction formats.

Comparing Auction Outcomes

From an economic perspective, there are three key dimensions in which auction formats can be compared, based on (i) revenues to the seller, (ii) rents to buyers, and (iii) efficiency. We begin by comparing the first-price and second-price auction outcome under the independent private value framework with risk-neutral buyers. Then we explore how the auction comparison looks like when we depart from this set of assumptions.

A central result in the auction literature is the equivalence theorem expressed in terms of expected revenues, expected utilities, and efficiency; see Vickrey (1961), Riley and Samuelson (1981), and Myerson (1981) in increasing generality. We shall consider the comparison of first-price and second-price auctions only. The result has been extended to wider class of auctions. The following theorem is based on the dominant strategy equilibrium in the second-price auction characterized in Theorem 1. We shall return to the pooling equilibrium in second-price auctions later on in the section on empirics of auctions.

Theorem 4 (Equivalence) Consider the symmetric independent private value framework with risk-neutral buyers. The following properties hold:

- (a) The expected revenue to the seller is the same in the first-price and second-price auction.
- (b) The item goes to the buyer who values it the most in the first-price and second-price auction.
- (c) The expected utility to the buyer is the same in the first-price and second-price auction.

Proof (a) The proof is an immediate consequence of our earlier theorems. Theorem 2 shows that in a first-price auction, buyers bid the expected second highest valuation conditional on being the high valuation. Furthermore, the high-valuation bidder wins. The expected price paid equals the expected second highest valuation. Consider next the sincere bidding equilibrium in the second-price auction described in Theorem 1. The high-valuation bidder wins at a price equal to the second highest valuation. The price paid equals the realization of the second highest

valuation. In expectations, the realization will equal the expected value.

(b) In the standard auction formats, the high-valuation bidder wins the item. The reason is that the bid functions are identical across bidders and strict monotone.

(c) From part (b), the allocation is the same in both a second-price auction and a first-price auction. In both cases, the high-valuation buyer wins. Part (a) shows that the expected revenues are the same. Therefore, expected utility must be the same.

Notice that parts (a) and (c) are in terms of expectations. Revenue, or utility, realizations need not be the same as the price paid may differ across auction formats.

How robust is the above theorem to departures from the assumptions? We shall see that the result is very fragile. If any of the assumptions is modified, the equivalence result breaks down. We shall discuss some contributions in this literature.

Suppose buyers are risk averse, instead of risk neutral, while maintaining all other assumptions. In a second-price auction, the (weakly dominant strategy) equilibrium is not affected. It remains an equilibrium that bidders bid their value. The equilibrium construction and proof of Theorem 1 remain valid in this case. Next consider a first-price auction. Maskin and Riley (1984) show that the equilibrium changes. Bidders bid more aggressively. The intuition is that bidders face the following trade-off. Bidding higher increases the chances of winning but comes at a utility loss of the higher bid. The first effect is not affected by attitudes toward risk, but the second is valued less when bidders are risk averse instead of risk neutral. In terms of the revenue ranking, this means a seller is better off with a first-price auction, while buyers prefer the second-price auction; see Matthews (1987).

Suppose bidders are asymmetric, that is, bidders draw their valuation from distinct probability distribution functions, while maintaining all other assumptions. Maskin and Riley (2000) have shown that revenue (and utility) ranking can go either way.

There are parametric examples of valuation distributions in which the first-price auction does

better for the seller and examples in which it does worse. In terms of efficiency, the second-price auction equilibrium remains efficient, while the first-price auction equilibrium is no longer efficient.

Third, consider interdependent valuations. The classic paper analyzing bidding equilibria in this case, also permitting that signals are correlated, is Milgrom and Weber (1982). They show that with affiliated interdependent valuations, the English auction performs best in terms of revenues followed by the second-price and then the first-price auction.

Myerson (1981) characterizes the revenue-maximizing auction. This pioneering paper develops a new approach, in which the auction rule is the choice variable. Myerson's studies auction rules from a mechanism design perspective in which buyers announce their valuations and the mechanism determines the allocation and transfer payments. Myerson shows that with independent private values, both first- and second-price auctions are optimal, but not if bidders' private values are asymmetric or correlated.

So far, we have considered the choice of auction format. Within an auction format, the seller can fine-tune the auction outcome. The seller may decide a minimum bid level below which bids are rejected, whether to charge bidder participation fees, or how much information about the object to make available to bidders. The minimum bid level or reserve price is easily understood with a single bidder. In the absence of a reserve price, the bidder would acquire the item at a price of zero. With a reserve price, the seller can force the bidder to pay a positive price at the cost of sometimes not selling the item. The optimal reserve is achieved at the point where the marginal benefit of increasing the reserve price becomes zero. Myerson (1981) gives an intuitive interpretation for the optimal reserve price formula. It is the valuation where information revelation costs (or incentive costs) equal the benefits from information.

The precision of information available to bidders can be influenced, by either the auctioneer or the buyers. For example, in a used car auction, the seller can decide whether bidders may inspect or even test-drive the car prior to the auction. Similarly, on eBay, the seller can decide on the

informativeness of the item description. In oil auctions bidders may decide how much money to spend on geological studies. Persico (2000) considers an affiliated-values environment and shows that bidder incentives to acquire information differ across auctions with the marginal benefit of additional information being higher in a first-price than in a second-price auction which may overturn Milgrom and Weber (1982) revenue ranking result. Bergemann and Pesendorfer (2007) study the joint decision problem when the seller controls the information and the auction rule in a private value setting. They show that increased information has the benefit of enhancing efficiency but comes at an information revelation cost. The optimal seller's policy has to balance these two elements. With few bidders, providing little information is optimal, while when the number of bidders increases, the efficiency motives dominate the information revelation cost element.

This section has examined some optimal auction design questions. Next, we shall consider empirics of auctions.

Empirics of Auctions

Bid data are available for many auctions and allow researchers to study bidders' behavior. The empirical literature has focused on two central questions: first, how to measure or quantify the underlying informational distribution from bid data and, second, how to design and assess the optimal auction for a market at hand. The first question in terms of econometrics is about the identification and inference of parameters determining the distribution of information. The second question is motivated by the fragility of the revenue equivalence theorem. Which elements, bidder asymmetry, risk aversion, and common versus private values, are key drivers? An answer to the second will tell us what auction rule is best used in practice, a market design question.

This section illustrates some empirical issues based on hypothetical and stylized data set. We shall ignore bidder heterogeneity, auction heterogeneity, and covariates. We shall comment on

these extensions later on. We shall start with first-price auctions and then consider second-price (or English) auction data.

Before proceeding, let us raise one central issue for empirical work which concerns the well-known problem of selection bias. The theoretical analysis and equilibrium characterization for specific auction rules assumes a known number of "potential" bidders. In fact, not all of these potential bidders may submit a bid. For example, a reserve price or bidder participation cost may reduce bidder participation. Empirically, this poses a problem as we only observe the "actual" bidders and need to infer the set of "potential" bidders. Put differently, the number of observed bids may not be an accurate picture of the degree of competition. One way to proceed is to estimate the potential number of bidders by using the maximum number of observed bids across auctions or a subset of auctions. Such an extremum estimator has nice asymptotic properties. Another way is to model this selection explicitly; see Li and Zheng (2009) and Athey et al. (2011). We shall ignore this issue here in this exposition and assume that the actual number of bidders equals the potential number of bidders.

Empirics of First-Price Auctions

Consider the following assumptions about the data-generating process consisting of a cross section of first-price sealed-bid auctions.

Assumption The data-generating process for the bid data $(b_i^t)_{i=1}^N$ is equilibrium bidding for independent private-value first-price auctions in which each auction t , $t = 1, \dots, T$ had (i) an identical object and (ii) a fixed (and known) number of identical bidders N .

Equilibrium bidding means that bidders follow the unique bidding strategy characterized in Theorem 2. Independent private values mean that a bidder i 's valuation equals her signal, $v_i = \theta_i$, and is drawn identically and independently from a distribution F . The fixed number of bidders means that we do not have to worry about the distinction between "actual" and "potential" bidders.

The empirical question is how to estimate the cdf F from bid data. Distinct estimation methods exist. Our description shall focus on a popular and commonly used inference approach. This approach looks at the optimal bid choice vis-à-vis the empirical distribution of opponent bids.

Following Guerre et al. (2000), the problem can be formulated based on $H(b)$ the probability distribution function of bids b . Bidder i 's problem of finding a bid that maximizes the expected pay-off in a first-price auction under risk-neutrality, private values, and when bids are independently distributed, can be written as:

$$\max_{b_i} [v_i - b_i]H(b_i)^{N-1}.$$

The first-order condition is

$$-H(b_i)^{N-1} + [v_i - b_i](N-1)H(b_i)^{N-2}H'(b_i) = 0,$$

which can be rewritten to obtain an explicit expression for the valuation v_i

$$v_i = b_i + \frac{H(b_i)}{(N-1)H'(b_i)}. \quad (5)$$

Equation 5 is the inverse of the theoretical bid function characterized in the section entitled "First-Price Sealed-Bid Auction". It tells us which valuation rationalizes the observed bid. The right-hand side elements are the bid b_i , and the number of bidders N , the cdf H , and the pdf H' . The pdf and cdf can be estimated from the bid data $(b_i^t)_{i,t}$ by using a suitable estimator. For example, the cdf H can be consistently estimated with the empirical cdf $\hat{H}(x) = \frac{1}{TN} \sum_{i,t} 1(b_i^t \leq x)$ where $1(\cdot)$ is the indicator function. The indicator function equals one, if the argument is nonnegative, and zero, otherwise.

Guerre et al. (2000) advocate a nonparametric estimator in which first the cdf and pdf of bids are estimated by using kernel estimators and then in the second step, (pseudo) valuations and their distribution are inferred. In practice this approach has gained a lot of popularity. One reason is its simplicity. A second reason is that it extends readily in various directions including bidder heterogeneity, different informational distributions, different auction rules, multi-unit auctions, and

even sequential auctions. In practice, researchers tend to use parametric approaches to estimate the distributions to allow for covariates to enter; see Hong and Paarsch (2006).

Empirics of Second-Price Auctions

Suppose the data-generating process is a second-price auction instead of a first-price auction.

Assumption The data-generating process for the bid data $(b_i^t)_{i=1}^N$ is equilibrium bidding for independent private-value second-price auctions in which each auction t , $t = 1, \dots, T$ had (i) an identical object and (ii) a fixed number of identical bidders N .

We have seen earlier that there can be multiple equilibria in second-price auctions. One equilibrium, in which one bidder bids high and all other bidders bid low, which is called a "pooling" equilibrium, has the feature that multiple valuations rationalize a bid. If this equilibrium arises in the data, then it may not be possible to infer the details of the distribution of valuations. However, bounds on the support of valuations could be inferred.

On the other hand, if the dominant strategy equilibrium is played, in which a bid equals the value, then inference of the valuation distribution is straightforward. In this case, the distribution of valuations can be estimated by using the empirical cdf of bids

$$\hat{F}(x) = \frac{1}{TN} \sum_{i,t} 1(b_i^t \leq x).$$

The dominant strategy equilibrium allows the econometrician to readily infer the underlying valuations and distribution of valuations from the observed bids.

In practice the researcher may not know the type of equilibrium played. Furthermore, different equilibria may be played in the cross section of auctions. How to deal with inference in this case has been an ongoing research area not only in the empirical auction literature but in economics in general; see Tamer (2003).

Winner's Curse

In addition to using field data, economists also use laboratory experiments to study market outcomes.

Bazerman and Samuelson (1983) conducted a first-price sealed-bid auction experiment with MBA students at Boston University. The number of student bidders varied between 4 and 26 across classes, and in each class, a jar of 800 pennies was offered for sale. The value of the jar was unknown to students. Students were asked to provide their bid and best estimate for the jar value. The average value estimate equaled \$5.13 which is \$2.87 below the true value. The average winning bid was \$10.01 which amounts to an average loss of \$2.01 with losses occurring in over half of all the auctions. The evidence suggests a “winner’s curse.”

The curse can be explained by using a (pure) common value environment in which bidder i ’s signal θ_i is bidder i ’s unbiased estimate of the value v , $E[\theta_i | v] = v$ (A framework that satisfies this assumption is when signals are noisy estimates of the true jar value, $\theta_i = v + \varepsilon_i$ with $E\varepsilon_i = 0$ iid.). If bidders use a monotone increasing bidding strategy, then the auction winner will be the bidder with the high signal, $\max_i \{\theta_i\}$. The curse arises as the high-signal bidder in fact overestimated the true value, $E[\max_i \{\theta_i\} | v] > \max_i E[\theta_i | v] = v$. The result follows since max is a convex function and from Jensen’s inequality, which says that for a convex function f , it must be that $Ef(\theta) > f(E\theta)$. Thus, winning confers bad news in the sense that the bidder learns ex post that she was the high-signal bidder.

How can we interpret the winner’s curse in terms of equilibrium bidding? The section entitled “First-Price Sealed-Bid Auction” has shown that bidders will shade their bid down in a private value setting. Now, with common values, bidders will shade their bid down even further in anticipation of the bad news effect. In particular, bidders will calculate the value of winning based on having the high signal for the object, $E[v | \theta_i, \theta_i \geq \theta_j \text{ for all } j]$.

Evidence from field data about a winner’s curse is mixed. Porter (1995) summarizes his joint work with Hendricks on field data for common value auctions. They analyze the US offshore mineral rights program in which the right to extract oil of the US coast has been awarded in form of a first-price sealed-bid auction since the 1950s.

Table 1 reports selected summary statistics for newly explored areas, Wildcat tracts, see Porter (1995, Table II). All dollar figures are in million of 1972 dollars. Standard deviations are in parenthesis. To the extent that the value of the oil field is the same for all bidders, say because of competitive raw oil prices, the informational environment can be viewed as common values.

Interestingly there is substantial uncertainty about the price of these oil fields. The amount overpaid by the winning bidder, $(b_{(1)} - b_{(2)})/b_{(2)}$, or “money left on the table,” averages to 44% or about 2.67 million of the average price of 6.07 million dollars. The uncertainty does remain substantial even when the number of bidders equals ten or more and amounts to 30% of the final price paid.

A winner’s curse implies a positive correlation between the price paid and the number of bids. Indeed, the data exhibit a positive correlation, as is evident in row 1. With one bidder, the average price paid, $b_{(1)}$, equals 1.50 million dollars, while with ten or more bidders, the price paid increases to 21.8 million. The price paid increases as the number of bidders increases. This positive correlation can alternatively be attributable to endogenous bidder participation decisions. For example, high-valued tracts may attract more bidders. To consider this alternative hypothesis, Table 1 reports what fraction of tracts were drilled and also productive (oil was found). Conditional on being drilled and oil is found, the last row reports the discounted revenues (ignoring drilling costs). Table 1 shows that tracts of higher value attract more bidders.

Table 2 reports a crude assessment of net returns calculated from Table 1. Net returns are defined as discounted revenues times the probability of drilling times the probability of being productive minus the winning bid. This calculation ignores drilling costs. Table 2 shows that net returns are positive throughout. There is no evidence that winning bidders make a loss in these oil field auctions. The evidence appears to reject the winner’s curse.

Investigating bidding strategies further, Porter reports an interesting (ex post) best response test which is a precursor to the (interim) best response formula developed by Guerre et al. (2000).

Auctions, Table 1 Summary statistics for Wildcat tracts

No of bidders	1	2	3	4	5–6	7–9	10–18	Total
$b_{(1)}$	1.50	2.76	4.17	5.62	7.9	14.2	21.8	6.07
	(0.1)	(0.2)	(0.4)	(0.5)	(0.6)	(1.2)	(1.4)	(0.2)
$(b_{(1)} - b_{(2)})/b_{(2)}$	–	0.55	0.49	0.46	0.39	0.34	0.30	0.44
Fraction drilled	0.61	0.74	0.86	0.85	0.91	0.90	0.99	0.78
Fraction productive	0.41	0.47	0.47	0.51	0.49	0.63	0.69	0.50
Discounted revenues	13.5	15.5	19.5	25.1	26.2	28.8	33.4	22.5
	(2.0)	(2.1)	(2.5)	(4.1)	(3.9)	(3.3)	(5.1)	(1.3)

Auctions, Table 2 Net returns

No of bidders	1	2	3	4	5–6	7–9	10–18	Total
Net returns	1.9	2.6	3.7	5.3	3.8	2.1	1.0	2.7

Consider the set of tracts on which a given firm submits bids. Assume that the bids of rival firms and ex post returns are held fixed. Suppose the vector of bids submitted by the firm in question is varied proportionally. If all of the firm's bids are increased, it will win more tracts but earn less per tract. In that way the optimal bid proportion that maximizes ex post returns can be calculated. Porter reports that few firms did not behave optimally and overbid. The calculation is ex post and does not take into account the uncertainty bidders were facing at the time of bid submission.

So far our analysis assumed that bidders behave competitively. Next, we shall describe some issues when bidders collude.

Collusive Bidding

Collusive bidding is illegal in many settings, but the temptation exists which has led to many bid rigging cases pursued by antitrust authorities. Exceptions include joint bidding in OCS auctions, which is legal at least among some bidders as described in Porter (1995), and subcontracting, which is legal in some procurement auctions. We shall discuss some issues relating to collusion among bidders. Collusion may also arise between the auctioneer and one or more bidders, but we shall leave that aside.

A bidding ring has to designate a cartel bidder. Suppose all bidders collude, so that the bidding ring is all-inclusive. Ideally, the ring would like to send only one bidder to the seller's auction and ask all other ring members to refrain from bidding or to submit phony bids. How can the ring determine their designated cartel bidder? Graham and Marshall (1987) point out that a pre-auction knockout can achieve this. Suppose the colluding ring holds an English auction prior to the seller's auction and shares the proceeds periodically in equal shares. The winner of the pre-auction knockout gets the right to be the only (serious) bidder at the seller's auction and pays the difference between the second highest bid and the seller's reserve price, $b_2 - R$. Theorem 1 shows that there exists an equilibrium in which bidders bid their value. Thus, the high-value bidder wins the knock-out auction and pays the second highest bid minus the seller's reserve price (provided the value is above the reserve price).

Observe that the collusive outcome achieves an efficient allocation. The high-value bidder wins. The rent distribution is now shifted relative to the competitive outcome, with a larger share of the rent going to bidders instead of the seller. What could the auctioneer do in response? Well, anticipating that a ring is in operation, the seller can increase the reserve price. The optimal level would be based on the expectation that the seller

faces a single cartel bidder with valuation being the high valuation from all ring.

The pre-auction knockout requires periodic division of spoils. Such side payments may leave a paper trail which increases the risk of detection by antitrust authorities. To minimize the risk of detection, the ring may refrain from using side payments all together. Instead the ring may use some other scheme or mechanism. Let $q_i(\hat{v})$ denote the probability that bidder i is the designated cartel bidder when ring members announce a valuation profile $\hat{v} = (\hat{v}_1, \dots, \hat{v}_N)$ to the ring mechanism. For bidders to tell the truth, it has to be that the expected probability of being the designated cartel bidder is independent of the own announcement $\hat{v}_i$. Otherwise, bidder i would announce the valuation that achieves the high probability. In turn this implies that the ring mechanism must assign items irrespective of the valuation. One scheme that achieves this is the “phases of the moon” allocation scheme, as used by companies involved in the great electrical conspiracy during the 1960s. Notice though that such a scheme is not efficient as not necessarily the bidder with the high valuation is selected. Thus, the collusive spoils will not be as high as when side payments are available.

The empirical literature on collusion in auctions is small. Two questions have been the focus: first, how can collusive behavior be detected based on bid data and, second, how cartels behave in practice. Porter and Zona (1993) propose a statistical test procedure to determine whether a subset of bidders colluded or not. The test is applied to highway procurement auctions. Pesendorfer (2000) shows that cartels may adopt distinct collusive schemes in practice. Pesendorfer studies school milk auctions and finds that in one regional market, Florida, cartel firms appear to use side payments, while in another regional market, Texas, cartel firms refrained from using side payments.

Concluding Remarks

Since the seminal papers by Vickrey (1961) and Milgrom and Weber (1982), research on auctions has created a large body of literature. This entry

examined well-established results including competitive bidding behavior in standard auction formats, revenue and utility comparison across auction formats, empirics of auctions, winner’s curse, and collusion. The setup involved a one-shot single-unit auction using a standard auction format.

Richer settings involving more goods, sold sequentially or simultaneously, and/or more involved market rules have already received attention in the economic literature and will attract more interest in the future. Market and mechanism design has become a successful area in economics from a theoretical, empirical, and practical perspective. It is to be expected that this will continue to be a fruitful research area in the future.

Bibliography

- Athey S, Levin J, Seira E (2011) Comparing open and sealed bid auctions: evidence from timber auctions. *Q J Econ* 126:207–257
- Bazerman MH, Samuelson WF (1983) I won the auction but don’t want the prize. *J Confl Resolut* 27(4):618–634
- Bergemann D, Pesendorfer M (2007) Information structures in optimal auctions. *J Econ Theory* 137:580–609
- Bergemann D, Brooks B, Morris S (2017) First-price auctions with general information structures: implications for bidding and revenue. *Econometrica* 85(1):107–143
- Cassady R (1967) *Auctions and Auctioneering*. University of California Press, Berkeley
- Graham DA, Marshall RC (1987) Collusive bidder behavior at single-object second-price and English auctions. *J Polit Econ* 95(6):1217–1239
- Guerre E, Perrigne I, Vuong Q (2000) Optimal nonparametric estimation of first-price auctions. *Econometrica* 68(3):525–574
- Hong H, Paarsch HH (2006) *An introduction to the structural econometrics of auction data*. MIT Press, Cambridge
- Hortacsu A, McAdams D (2016) Empirical work on auctions of multiple objects. *J Econ Lit* (forthcoming)
- Klemperer P (1999) *Auction theory: a guide to the literature*. *J Econ Surv* 13(3):227–286
- Krishna V (2002) *Auction theory*. Academic, San Diego
- Li T, Zheng X (2009) Entry and competition effects in first-price auctions: theory and evidence from procurement auctions. *Rev Econ Stud* 76(4):1397–1429
- Maskin E, Riley J (1984) Optimal auctions with risk averse buyers. *Econometrica* 52(6):1473–1518
- Maskin E, Riley J (2000) Asymmetric auctions. *Rev Econ Stud* 67:413–438

- Maskin E, Riley J (2003) Uniqueness of equilibrium in sealed high-bid auctions. *Games Econ Behav* 45:395–409
- Matthews S (1987) Comparing auctions for risk averse buyers: a buyer's point of view. *Econometrica* 55(3):633–646
- McAfee P, McMillan J (1987) Auctions and bidding. *J Econ Lit* 25:699–738
- Milgrom PR, Weber RJ (1982) A theory of auctions and competitive bidding. *Econometrica* 50:1089–1122
- Myerson RB (1981) Optimal auction design. *Math Oper Res* 6:58–73
- Persico N (2000) Information acquisition in auctions. *Econometrica* 68(1):135–148
- Pesendorfer M (2000) A study of collusion in first-price auctions. *Rev Econ Stud* 67(3):381–411
- Porter RH (1995) The role of information in U.-S. offshore oil and gas lease auctions. *Econometrica* 63(1):1–27
- Porter RH, Zona JD (1993) Detection of bid rigging in procurement auctions. *J Polit Econ* 101(3): 518–538
- Riley JG, Samuelson WF (1981) Optimal auctions. *Am Econ Rev* 71:381–392
- Tamer E (2003) Incomplete simultaneous discrete response model with multiple equilibria. *Rev Econ Stud* 70(1):147–165
- Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. *J Financ* 16(1):8–37
- Wilson R (1977) A bidding model of perfect competition. *Rev Econ Stud* 44:511–518

Implementation Theory

Luis C. Corchón
Departamento de Economía, Universidad
Carlos III, Madrid, Spain

Article Outline

Glossary
Definition
Introduction
Brief History of Implementation Theory
The Main Concepts
The Main Insights
Unsolved Issues and Further Research
Answers to the Questions
Bibliography

Glossary

Equilibrium concept A mapping (or a collection of them) from the set of states of the world into allocations yielded by equilibrium messages. This equilibrium is a game-theoretical notion of how agents behave, e.g., Nash equilibrium, Bayesian equilibrium, dominant strategies, etc.

Implementable social choice rule in an equilibrium concept (e.g., Nash equilibrium) A social choice rule is implementable in an equilibrium concept (e.g., Nash equilibrium) if there is a mechanism such that for each state of the world, the allocations prescribed by the social choice rule and those yielded by the equilibrium concept coincide.

Mechanism A list of message spaces and an outcome function mapping messages into allocations. It represents the communication and decision aspects of the organization.

Social choice rule A correspondence mapping the set of states of the world in the set of allocations. It represents the social objectives that the society or its representatives want to achieve.

State of the world Description of all information possessed by all agents.

Type of an agent All the information possessed by this agent. It may refer to the preferences of this agent and/or to the knowledge of this agent of the preferences of other agents.

Definition

Implementation theory studies which social objectives (i.e., social choice rules) are compatible with the incentives of the agents (i.e., are implementable). In other words, it is the systematic study of the social goals that can be achieved when agents behave strategically.

Introduction

Dear colleague;

I wrote this survey with you in mind. You are an economist doing research who would like to know why implementation is important. And by this I do not mean why some people won the Nobel Prize working in this area. I mean, what are the deep insights found by implementation theory and what applications are delivered by these tools. I propose a simple game: try to answer the following questions. If you cannot answer them, but you think they are important, read the survey. At the end of this survey, I will give you the answers. I will also tell you why I like implementation theory so much!

1. Why are agents price-takers? Is price-taking possible in economies with a finite number of agents?
2. Suppose two firms wish to merge. They claim that the merger will bring large cost reductions but some people fear that the firms just want to avoid competition. What would be your advice?
3. How should a monopoly be regulated when regulators do not know the cost function or the demand function of the monopolist?

4. How should it be determined whether or not a public facility – a road, a bridge, and a stadium – should be constructed and who should pay for it?
5. Is justice possible in this world? Can we reconcile justice and self-interest?
6. Can an uninformed planner achieve better allocations than those produced by completely informed agents in an unregulated market?
7. In competitive ice skating, the highest and lowest marks awarded by judges are discarded and the remaining are averaged. Do you think that this procedure eliminates incentives to manipulate votes?
8. What kind of policies would you advocate to fight global warming?

The answers to these questions are found in section “Answers to Questions.” The rest of this paper goes as follows. Section “[Brief History of Implementation Theory](#)” is a historical introduction that can be skipped. Section “[The Main Concepts](#)” explains the basic model. Section “[The Main Insights](#)” explains the main results. Section “[Unsolved Issues and Further Research](#)” offers some thoughts about the future direction of the topic.

Brief History of Implementation Theory

From, at least Adam Smith on, we have assumed that agents are motivated by self-interest. We also assumed that agents interact in a market economy where prices match supply and demand. This tradition crystallized in the Arrow-Debreu-McKenzie model of general equilibrium in the 1950s. But it was quickly discovered that this model had important pitfalls other than focusing on a narrow class of economic systems: On the one hand, an extra agent was needed to set prices, the auctioneer. On the other hand, agents follow rules, i.e., to take prices as given, which are not necessarily consistent with self-interest. An identical question had arisen earlier when Taylor (1929) and Lange (1936–1937), following Barone (1908), proposed a market socialism, where socialist managers maximize profits: Why would socialist managers choose output in the way

prescribed to them (or who will provide and preserve capital in a system where the private property of such items is forbidden)? Samuelson (1954) voiced identical concern about the Lindahl solution to allocate public goods: “It is in the selfish interest of each person to give false signals.” This concern gave rise later on to the golden rule of incentives – as stated by Roger Myerson (1985): “An organization must give its members the correct incentives to share information and act appropriately.” Earlier, it had aroused the interest of Leonid Hurwicz, the father of implementation theory, in economic systems other than the market. In any case, it was clear that an important ingredient was missing in the theory of economic systems. This element was that not all information needed for resource allocation was transmitted by prices: Some vital items have to be transmitted by agents.

Several proposals arose to fill the gap: on the one hand, models of markets under asymmetric information, Vickrey (1961), Akerlof (1970), Spence (1973), and Rothchild and Stiglitz (1976), and on the other hand, models of public intervention, like optimal taxation, Mirless (1971), and mechanisms for allocating public goods, Clarke (1971) and Groves (1973), with the so-called principal-agent models somewhere in the middle. The keyword was “truthful revelation” or “incentive compatibility”: Truthful revelation of information must be an equilibrium strategy, either a dominant strategy, as in Clarke and Groves, or a Bayesian equilibrium as in Arrow (1977) and D’Aspremont and Gerard-Varet (1979). A motivation for this procedure was provided by the “revelation principle,” Gibbard (1973), Myerson (1979), Dasgupta et al. (1979), and Harris and Townsend (1981): If a mechanism yields certain allocations in equilibrium, telling the truth about one’s characteristics must be an equilibrium as well (however, telling the truth may not be an equilibrium in the original mechanism you might have to use an equivalent direct mechanism). This result is of utmost importance and it will be thoroughly considered in section “[The Main Concepts](#).” However, it was somehow misread as “there is no loss of generality in focussing on incentive compatibility.” But what the revelation principle asserts is that truthful revelation is one of

the, possibly, many equilibria. It does not say that truthful revelation is *the only* equilibrium. As we will see in some cases, it is a particularly unsatisfactory way of selecting equilibria.

The paper by Hurwicz (1959), popularized by Reiter (1977), presented a formal structure for the study of economic mechanisms which has been followed by all subsequent papers. Maskin (1999), whose first version circulated in 1977, is credited as the first paper where the problem of multiple equilibria was addressed as a part of the model and not as an afterthought; see the report of the Nobel Prize Committee (2007). Maskin studied implementation in Nash equilibrium (see “Glossary”). Later his results were generalized to Bayesian equilibrium by Postlewaite and Schmeidler (1986) and Palfrey and Srivastava (1987, 1989).

Finally, Moulin (1979) studied dominance solvability and Moore and Repullo (1988) subgame perfect equilibrium. The century closed with several characterizations on what can be implemented in other equilibrium concepts: Moore and Repullo (1990) in Nash equilibrium, Palfrey and Srivastava (1991) in undominated Nash equilibrium, Jackson (1991) in Bayesian equilibrium, Dutta and Sen (1991a) in strong equilibrium, and Sjöström (1993) in trembling hand equilibria. With all these papers in mind, the basic aspects of implementation theory are now well understood.

The interested reader may complement the previous account with the surveys by Maskin and Sjöström (2002) and Serrano (2004) which cover the basic results and by Baliga and Sjöström (2007) for new developments including experiments. See also Maskin (1985), Moore (1992), Corchón (1996), Jackson (2001), and Palfrey (2002). Several important applications of implementation theory are not surveyed here: auctions, see Krishna (2002); contract theory, see Lafont and Martimort (2001); matching, see Roth (2008); and moral hazard, see Ma et al. (1988).

The Main Concepts

We divide this section into four subsections: The first describes the environment, the second deals

with social objectives, the third revolves around the notion of a mechanism, and the last defines the equilibrium concepts that we will use here.

The Environment

Let $I = \{1, \dots, n\}$ be the set of agents. Let θ_i be the type of i . This includes all the information in the hands of i . Let Θ_i be agent i 's type set. The set $\Theta \subset \prod_{i=1}^n \Theta_i$ is the set of states of the world. For each state of the world, we have a feasible set $A(\theta)$ and a preference profile $R(\theta) = (R_1(\theta), \dots, R_n(\theta))$. $R_i(\theta)$ is a complete, reflexive, and transitive binary relation on $A(\theta)$. $I_i(\theta)$ denotes the corresponding indifference relation. Set $A \equiv \bigcup_{\theta \in \Theta} A(\theta)$. Let $a = (a_1, a_2, \dots, a_n) \in A$ be an allocation, also written (a_i, a_{-i}) , where $a_{-i} \equiv (a_1, a_2, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$.

The standard model of an exchange economy is a special case of this model: θ is an economy. $X_i(\theta) \subset \mathfrak{R}^k$ is the consumption set of i . $w_i(\theta) \in \text{int} X_i(\theta)$ are the endowments in the hands of i . The preferences of i are defined on $X_i(\theta)$. The set of allocations $A(\theta)$ is defined as

$$A(\theta) = \left\{ a \left| \sum_{i=1}^n (a_{ij} + w_{ij}(\theta)) \leq 0, j = 1, 2, \dots, k, \right. \right. \\ \left. \left. (a_{i1}, a_{i2}, \dots, a_{ik}) \in X_i(\theta), \forall i \in I. \right\}.$$

A special case of an exchange economy is bilateral trading: Here there are two agents, the seller and the buyer. The seller has a unit of an indivisible good and both agents are endowed with an infinitely divisible good (“money”). Preferences are representable by linear utility functions. The type of each agent, also called her valuation, is the marginal rate of substitution between both goods. Finally, the set of types is a closed interval of the real line.

Another example is the social choice model where the set of states of the world is the Cartesian product of individual type sets, $\Theta = \prod_{i=1}^n \Theta_i$. The set of feasible allocations is constant. The preferences of each agent only depend on her type, for all $\theta \in \Theta$, $R_i(\theta) = R_i(\theta_i)$ all $i \in I$.

The model of public goods is a hybrid of the social choice and the exchange economy models.

For a subset of goods, say $1, 2, \dots, l$, agents receive the same bundle (these are the public goods). For goods $l+1, \dots, k$, agents can consume possibly different bundles.

Social Objectives

Implementation begins by asking what allocations we want to achieve. In this sense, implementation theory reverses the usual procedure, namely, fix a mechanism and see what the outcomes are. The theory is rather agnostic as to who is behind *we*: It could be a democratic society, it could be a dictator, it could be a benevolent planner, etc. Formally, a correspondence $F : \Theta \rightarrow A$ such that $F(\theta) \subseteq A(\theta)$ for all $\theta \in \Theta$ will be called a *social choice rule* (SCR). Under risk or uncertainty, allocations are state-dependent (recall the concept of contingent commodities in general equilibrium). Thus, an allocation is a single-valued function $f : \Theta \rightarrow A$. The notion of a SCR is replaced by that of a *social choice set* (SCS) defined as a collection of functions mapping Θ into A . Examples of SCR are the Pareto rule, which maps every state into the set of Pareto efficient allocations for this state, the Walrasian SCR which maps every economy in the set of allocations that are a Walrasian equilibrium for this economy, etc.

If states of the world were contractible, i.e., if they could be written in an enforceable contract specifying the allocations in each state, SCR or SCS would be directly achieved, assuming that those not complying could be punished harshly. Unfortunately, states of the world are a description of preferences and productive capabilities, being those difficult to describe and therefore easy to manipulate. Thus, we have to find another method to reach the desired allocations.

Mechanisms

If the information necessary to judge the desirability of allocations is in the hands of agents, it seems that the only way of retrieving this information is by asking them. But, of course, agents cannot be trusted to reveal truthfully their information because they might lose by doing so. Thus, the owner of a defective car will think twice about revealing the true state of the car if the price of

defective cars is less than the price of reliable cars. But perhaps we may design ways in which the messages sent by different agents are checked one against the other. We may also design ways in which agents send information by indirect means, say by raising flags, making gestures, and so on and so forth. This is the idea behind the concept of a mechanism (also called a game form).

Formally, a *mechanism* is a pair (M, g) where $M \equiv \prod_i^n M_i$ is the *message space* and $g : M \rightarrow A$ is the outcome function. M_i denotes agent i 's *message space with typical element m_i* . In some cases, i.e., when goods are indivisible, the outcome function maps M into the set of lotteries on A , denoted by $\mathcal{L}A$. In this case, the outcome function yields the probability of obtaining an object. Let $m = (m_1, \dots, m_n) \in M$, be a list of messages, also written (m_i, m_{-i}) where m_{-i} is a list of all messages except those sent by i .

Another interpretation of a mechanism, more in tune with decentralized systems, is that messages describe contracts among agents, and the outcome function is a legal system that converts contracts into allocations.

If feasible sets are state-dependent, we have a problem: Suppose that at θ we want to achieve allocation $a \in A(\theta)$. So there must be a message, say m such that $g(m) = a$. But what if there is another state, say θ' for which $a \notin A(\theta')$? In this case, $g(m) \notin A(\theta')$. In other words, since mechanisms are not state-dependent, they may yield unfeasible allocations. We will postpone the discussion of this problem until section “[Unsolved Issues and Further Research](#).” For the time being, let us assume that feasible sets are not state-dependent.

Equilibrium

Since the messages sent by agents are tied to their incentives, it is clear that we have to use an equilibrium concept borrowed from game theory. Thus, given $\theta \in \Theta$, a mechanism (M, g) induces a game in normal form (M, g, θ) . There are many “solutions” to what would constitute an equilibrium. Let us begin by considering the notion of a Nash equilibrium.

Definition 1 A message profile $m^* \in M$ is a Nash equilibrium for (M, g, θ) if, for all $i \in I$ $g(m^*)R_i(\theta)g(m_i, m_{-i}^*)$ for all $m_i \in M_i$.

Let $NE(M, g, \theta)$ be the set of allocations yielded by all Nash equilibria of (M, g, θ) . We now ask, given a SCR, what mechanism, if any, would produce outcomes identical to the SCR? In this sense, the mechanism is the variable of our analysis, i.e., the mechanism “solves” the equation $NE(M, g, \theta) = F(\theta)$, for all $\theta \in \Theta$. Formally,

Definition 2 The SCR F is implementable in Nash equilibrium if there is a mechanism (M, g) such that, for all $\theta \in \Theta$, $NE(M, g, \theta) \neq \emptyset$ and:

1. $F(\theta) \subseteq NE(M, g, \theta)$.
2. $NE(M, g, \theta) \subseteq F(\theta)$.

The previous concept can be easily generalized. Given a mechanism (M, g) , an equilibrium concept is a mapping, say $E_{(M, g)} : \Theta \rightarrow A$ such that $E_{(M, g)}(\theta) \subseteq A(\theta)$ for all $\theta \in \Theta$. For instance, $E_{(M, g)}(\theta)$ may be the set of allocations arising from dominant strategy profiles in θ when the mechanism (M, g) is in place. The notion of implementation in an equilibrium concept easily follows. See Thomson (1996) for a discussion of other concepts of implementation.

The problem is that some equilibrium concepts cannot be written in the way we just described because the actions to be taken in state, say θ' , depend on preferences in states other than θ' . To see this, suppose that agents attach a vector of probabilities to each possible type of the other agents, Harsanyi (1967/1968). Denote by $q(\theta_{-i}/\theta_i)$ the vector of probabilities attached by i that other agents have types given that she is of type θ_i . For simplicity, assume that it is a strictly positive vector. Suppose that preferences are representable by a von Neumann-Morgenstern utility index $V_i(a, \theta)$. In this framework (as first noticed by Vickrey (1961)), a strategy for i , denoted by s_i , is no longer a message but a function from the set of types of i in the set of messages of i , namely, $s_i : \Theta \rightarrow M_i$. A strategy profile, s , is a collection of strategies, one for each agent, $s =$

$(s_1, \dots, s_n)$ also written as (s_i, s_{-i}) . For simplicity, the next definition assumes that type sets are finite.

Definition 3 A Bayesian equilibrium (BE) for $(M, g, R(\cdot))$ is a s^* such that for all i , $\theta \in \Theta$, and $m_i \in M_i$,

$$\begin{aligned} & \sum_{\theta_{-i} \in \Theta_{-i}} q(\theta_{-i} | \theta_i) V_i(g(s^*(\theta)), \theta) \\ & \geq \sum_{\theta_{-i} \in \Theta_{-i}} q(\theta_{-i} | \theta_i) V_i(g(m_i, s_{-i}^*(\theta_{-i})), \theta) \end{aligned}$$

Thus, an equilibrium concept – given a mechanism – is a collection of functions, denoted by $H_{(M, g)}$, such that for all $h_{(M, g)} \in H_{(M, g)} : \Theta \rightarrow A$. Finally, the definition of implementable SCS in BE follows.

Definition 4 The mechanism (M, g) implements a SCS F in BE if:

1. For any BE s , there exists $x \in F(\theta)$, such that $g(s(\theta)) = x(\theta)$ for all $\theta \in \Theta$.
2. For any $x \in F$, there is a BE s such that $g(s(\theta)) = x(\theta)$ for all $\theta \in \Theta$.

Looking at our definitions of an implementable SCR or SCS, we see that the first requirement is that all equilibria yield “good” allocations. The second requirement is that given an allocation to be implemented, there is an equilibrium “sustaining” this allocation. These two requirements bear some resemblance to the two fundamental theorems of welfare economics, namely, that competitive equilibrium is efficient and that any efficient allocation can be achieved as a competitive equilibrium with the appropriate endowment redistribution. Notice that endowment redistribution is not used in the definition of implementation.

The Main Insights

We group our results here under three headings: “The Revelation Principle and Its Consequences,” “Monotonicity and How to Avoid It,” and “The Limits of Design.” We will discuss them each in turn.

The Revelation Principle and Its Consequences

The definition of a mechanism is extremely abstract. No conditions have been imposed on what might constitute a message space or an outcome function. And since implementation theory considers the mechanism the variable to be found, this is an unhappy situation: We are asked to find something whose characteristics we do not know! Fortunately, the revelation principle comes to the rescue by stating a necessary condition for implementation: If a single-valued SCR, which we will call a social choice function (SCF), is implementable, there is a *revelation mechanism* for which telling the truth is an equilibrium. A *revelation mechanism* (associated with a SCF) is a mechanism in which the message space for each agent is her set of types and the outcome function is the SCF. We say that a SCF is truthfully implementable or incentive compatible if truth-telling is a Bayesian equilibrium (or a dominant strategy) of the direct mechanism associated with it. The following result formally states the revelation principle:

Theorem 1 *If f is a Bayesian (resp. dominant strategy) implementable SCF, f is incentive compatible.*

Proof Let f be Bayesian implementable. Therefore, there exists a mechanism (M, g) and a Bayesian equilibrium s^* such that $g(s^*(\theta)) = f(\theta)$ for every $\theta \in \Theta$. Since $s^*(-)$ is a Bayesian equilibrium, $\forall \theta \in \Theta, \forall m_i \in M$

$$\begin{aligned} & \sum_{\theta_{-i} \in \Theta_{-i}} q(\theta_{-i} | \theta) V_i(g(s^*(\theta)), \theta) \\ & \geq \sum_{\theta_{-i} \in \Theta_{-i}} q(\theta_{-i} | \theta_i) V_i(g(m_i, s_{-i}^*(\theta_{-i})), \theta), \end{aligned}$$

which implies that $\forall \theta'_i \in \Theta_i, \forall \theta_{-i} \in \Theta_{-i}$,

$$\begin{aligned} & \sum_{\theta_{-i} \in \Theta_{-i}} q(\theta_{-i} | \theta_i) V_i(f(\theta), \theta) \\ & \geq \sum_{\theta_{-i} \in \Theta_{-i}} q(\theta_{-i} | \theta_i) V_i(f(\theta'_i, \theta_{-i}), \theta). \end{aligned}$$

The proof for the case of dominant strategies is identical.

Theorem 1 (T. 1 in the sequel) can be explained in terms of a mediator, i.e., somebody to whom you say “who you are” and who chooses the strategy that maximizes your payoffs on your behalf. Would you try to fool such a person? If you do so, you are fooling yourself because the mediator would choose a strategy that is not the best for you. Thus, the best thing for you to do is to tell the truth (providing an unexpected backing to the aphorism “honesty is the best policy!”).

Consider now the following results, due to Hurwicz (1972) (who proved it for the case of $n = 2$) and to Gibbard (1973) and Satterthwaite (1975), respectively:

Theorem 2 *In exchange economy environments, there is no SCF such that:*

1. *It is truthfully implementable in dominant strategies.*
2. *It selects individually rational allocations.*
3. *It selects efficient allocations.*
4. *Its domain includes all economies with convex and continuous preferences.*

Theorem 2' *In social choice environments, there is no SCF such that:*

1. *It is truthfully implementable in dominant strategies.*
2. *It is non-dictatorial.*
3. *Its range is A , with $\#A > 2$.*
4. *Its domain includes all possible preference profiles.*

It is clear that there are trivial SCF in which any three conditions in T. 2–2' are compatible. But T. 2–2' are very robust in the sense that they hold for smaller domains of economies (Barberá and Peleg 1990; Barberá et al. 1991), for weaker notions of individual rationality (Saijo 1991; Serizawa and Weymark 2003) and in public goods domains (Ledyard and Roberts 1974). Moreover, assuming quasi-linear utility functions, Hurwicz and Walker (1990), building on a previous paper by Walker, proved that the set of economies for which conditions 1–3 in T. 2 are incompatible is open and dense. Beviá and Corchón (1995) show that these conditions are

incompatible for *any* economy where utility functions are quasi-linear, strictly concave, and differentiable and fulfill a very mild regularity condition. These results show that Vickrey-Clarke-Groves mechanisms fail to achieve efficient allocations in general (Vickrey-Clarke-Groves mechanisms are revelation mechanisms that work in public good economies where utility functions are quasi-linear in “money.” The outcome function selects the level of public good that maximizes the sum of utilities announced by agents, and the money received by each individual is the sum of the utility functions announced by all other agents. For an exposition of these mechanisms, see Green and Laffont (1979). The most general domain in which they achieve efficient allocations is in Tian (1996)).

A proof of T. 2 can be found in Serizawa (2002). Simple proofs of T. 2' can be found in Barberá (1983), Benoit (2000), and Sen (2001).

T. 1 and 2–2' imply that there is no mechanism implementing an efficient and individually rational (resp. non-dictatorial) SCF in dominant strategies when the domain of the SCF is large enough. In other words, the revelation principle implies that the restriction to mechanisms where agents announce their own characteristic is not important when considering negative results. Thus, the revelation principle is an appropriate tool for producing negative results. But we will see that to rely entirely on this principle when trying to implement a SCF may yield disastrous results.

Choosing wisely the domain, Barberá et al. (2010, 2016) have shown that many salient allocation rules are not only incentive compatible but also group incentive compatible (i.e., truth is a dominant strategy for any coordinated deviation for any group of players) in a variety of allocation problems like voting, matching, fair division, cost sharing, house allocation, and auctions. The conditions used by these authors are an example that the conflict between incentive compatibility and efficiency is solvable as long as we are prepared to restrict the scope of applicability of the mechanism.

A natural question to ask is what happens with the above impossibility results when we weaken the requirement of implementation in dominant strategies to that of implementation in Bayesian

equilibrium. The following result, due to Myerson and Satterthwaite (1983), answers this question:

Theorem 2'' *In the bilateral trading environment, there is no SCF such that:*

1. *It is truthfully implementable in Bayesian equilibrium.*
2. *It selects individually rational allocations once agents learn their types.*
3. *It selects ex post efficient allocations.*
4. *Its domain includes all linear utility functions with independent types distributed with positive density, and the sets of types have a non-empty intersection.*

Proof (Sketch; see Krishna and Perry 1997 for details) By the revenue equivalence theorem (see Klemperer 1999, Appendix A), all mechanisms fulfilling conditions (2) and (3) above raise identical revenue. So it is sufficient to consider the Vickrey-Clarke-Groves which, as we remarked before, is not efficient.

Again the weakening of any condition in T. 2'' may produce positive results (Williams (1999), Table 1, presents an illuminating discussion of this issue). For instance, suppose seller valuations are 1 or 3, and buyer valuations are 0 or 2. The mechanism fixes the price at 1.5 and a sale occurs when the valuation of the buyer is larger than the valuation of the seller. This mechanism implements truthfully a SCF satisfying (2) and (3) above. Unfortunately, it does not work when valuations are drawn from a common interval with positive densities.

But unlike T. 2–2', there are robust examples of SCF truthfully implementable in Bayesian equilibrium when conditions (2) or (4) are relaxed. Also, inefficiency converges to zero very quickly when the number of agents increases (see Gresik and Satterthwaite 1989). This is because the equilibrium concept is now weaker and we are approaching a land where incentive compatibility has no bite, as we will see in T. 3 below.

First, d'Aspremont and Gerard-Varet (1975, 1979) and Arrow (1979) showed that conditions (1)–(3)–(4) are compatible with individual rationality *before* agents learn their types in the domain

of public goods with quasi-linear utility functions. They proposed the “expected externality mechanism” in which each agent is charged the expected externality she creates on the remaining players. Later on, Myerson (1981) and Makowski and Mezzetti (1993) presented incentive compatible SCF yielding ex post efficient and individually rational allocations in the domain of exchange economies with quasi-linear preferences and more than two buyers. In Myerson (1981), agents have correlated valuations. Buyers are charged even if they do not obtain the object or they may receive money and no object or even receive the object plus some money. Makowski and Mezzetti (1993) assume no correlation and that the highest possible valuation for a buyer is larger than the seller’s highest possible valuation. They consider a family of mechanisms, called second price auction with seller (SPAWS), in which the highest bidder obtains the object, the seller receives the first bid, and the winning buyer pays the second price. These mechanisms not only induce truthful behavior and yield ex post efficient and individually rational allocations: For any other mechanism with these properties, we can find a SPAWS mechanism yielding the same allocation.

Suppose now that information is *nonexclusive* in the sense that the type of each player can be inferred from the knowledge of all the other players’ type. Intuition suggests that in this case, incentive compatibility has no bite whatsoever (i.e., T2'' does not apply) since the behavior of each player can be “policed” by the remaining players. In order to prove this, we will concentrate on an extreme, but illuminating, case of non-exclusive information, namely, Nash equilibrium. In this framework, since information is complete, a direct mechanism is one where each agent announces a state of the world.

Consider the following assumption:

- (W) $\exists z \in A$ such that $\forall \theta \in \Theta, \forall a \in A, aR_i(\theta)z, \forall i \in I$.

This assumption will be called “universally worst outcome” because it postulates the existence of an allocation which is unanimously deemed as the worst. In an exchange economy,

this allocation would be zero consumption for everybody. Now we have the following result (Repullo 1986; Matsushima 1988):

Theorem 3 *If $n = 2$ and W holds, any SCF is truthfully implementable in Nash equilibrium. If $n > 2$, any SCF is truthfully implementable in Nash equilibrium.*

Proof When $n = 2$, consider the following outcome function: $g(\theta', \theta') = f(\theta') \forall \theta' \in \Theta, g(\theta', \theta'') = z$ for all $\theta' \neq \theta''$. Clearly, truth is an equilibrium. When $n > 2$, consider the following outcome function: If m is such that $n - 1$ agents announce state θ' , then $g(m) = f(\theta')$. Otherwise, $g(\cdot)$ is arbitrary. Clearly, truth is an equilibrium as well in this case.

The first thing to notice is the difference between the cases of two and more than two individuals. We will have more to say about this in the next section. The second is that the construction in Theorem 3 produces a large number of equilibria and that there seems to be no good reason for individuals to coordinate in the truthful equilibria.

For instance, suppose workers can be either fit or unfit. When a profit-maximizing firm asks its employees about their characteristics, and all workers are fit, a unanimous announcement such as “we are all unfit” is an equilibrium. If fit workers are required hard work and unfit workers are asked to light work, do you think it is reasonable that workers coordinate in the truthful equilibrium? A more elaborate example was produced by Postlewaite and Schmeidler (1986): There are three agents. The first agent has no information and agents 2 and 3 are perfectly informed. The ranking of agent 1 over alternatives is the opposite of agents 2 and 3 who share the same preferences. The SCF is the top alternative of agent 1 in each state. It is intuitively clear that besides the truthful equilibria, there is another untruthful equilibrium where both informed agents lie, and they are strictly better off than under truthful behavior. Again, coordination in the truthful equilibrium seems very unlikely. Thus, we have to recognize that we have a problem here. The next section will tell you how we can solve it.

Summing up, what do we learn from the results in this section?

1. When looking for an implementable SCF, a useful first test is whether this SCF yields incentives for the agents to tell the truth; see T. 1. But this test is incomplete because of the existence of equilibria other than the truthful one; see T. 3. These untruthful equilibria sometimes sound more plausible than the truthful one.
2. All impossibility theorems – T. 2–2'–2'' – have the same structure: truthful implementation, individual rationality/non-dictatorship, efficiency/large range of the SCR, and large domain. Usually, in social choice environments, conditions 2 and 3 are weaker than in economic environments but the condition on the domain is stronger.
3. The classic story of the market making possible efficient allocation of resources under private information has to be revised. Private information in many cases precludes the existence of *any* mechanism achieving efficient and individually rational allocations under informational decentralization; see T. 2–2'–2''.
4. The same remarks apply to naive applications of the Coase theorem where agents are supposed to achieve Pareto efficient allocations just because they have contractual freedom (ditto about bargaining theory). In the parlance of Coase, private information is an important transaction cost.
5. When mechanisms with adequate properties exist, like those proposed by Arrow, d'Aspremont and Gerard-Varet, Myerson, and Makowski and Mezzetti, they are not of the kind that we see in the streets. Careful design is needed. These mechanisms are tailored to specific assumptions on valuations; thus, their range of applicability may be limited.

Monotonicity and How to Avoid It

We have seen that equilibria other than the truthful one are likely to arise. We have also seen that these equilibria cannot be disregarded a priori. So we have to find a way of getting rid of equilibria that do not yield desirable allocations. Under dominant

strategies, clearly, if all preference orderings are strict, implementation and truthful implementation becomes identical; see Dasgupta et al. (1979), Corollary 4.1.4 (Laffont and Maskin 1982 present other conditions under which this result holds. See Repullo (1985) for the case where implementation and truthful implementation in dominant strategies do not coincide). For the ease of exposition, we consider next Nash equilibria.

It turns out that the key to this issue in the case of Nash equilibrium is the following monotonicity property, sometimes called Maskin monotonicity because Maskin (1977) established its central relevance to implementation:

- (M) A SCR F is monotonic if

$$\{a \in F(\theta), aR_i(\theta)b \rightarrow aR_i(\theta')b \quad \forall i \in I\} \\ \rightarrow a \in F(\theta').$$

Monotonicity says that if an allocation is chosen in state θ and this allocation does not fall in anybody's ranking in state θ' , this allocation must also be chosen in θ' . We will also speak of a "monotonic transformation of preferences at θ'' " when the requirement $aR_i(\theta)b \rightarrow aR_i(\theta')b \quad \forall i \in I$ is satisfied. This requirement simply says that the set of preferred allocations shrinks when we go from θ to θ' .

Monotonicity looks like a not unreasonable property, even though, as we will see in a moment, there are cases in which it is incompatible with other very desirable properties. In any case, the importance of monotonicity comes from the fact that it is a necessary condition for implementation in Nash equilibrium, as proved by Maskin (1977).

Theorem 4 *If a SCR is implementable in Nash equilibrium, it is monotonic.*

Proof If F is Nash implementable, there must be a mechanism (M, g) such that $\forall a \in F(\theta)$; there is a Nash equilibrium m^* , such that $g(m^*) = a$. Since $aR_i(\theta)b \rightarrow aR_i(\theta')b \quad \forall i \in I$, m^* is also a Nash equilibrium at θ' . Since F is implementable, $a \in F(\theta')$.

Let us now discuss the concept of monotonicity. First, the bad news. Popular concepts in voting, like plurality, Borda scoring, and majority rule, are not monotonic, neither is the Pareto correspondence; see Palfrey and Srivastava (1991), p. 484. Even the venerable Walrasian correspondence is not monotonic! The failure of the Pareto and the Walrasian SCR to be monotonic can be amended: If preferences are strictly increasing in all goods, the Pareto SCR is monotonic in economic environments. The *constrained Walrasian* SCR – in which consumers maximize with respect to the budget constraint and the availability of resources – is also monotonic. More serious is a result due to Hurwicz (1979) that uses two weak conditions on a SCR defined in the domain of exchange economies:

- (L) The domain of F contains all preferences representable by linear utility functions.
- (ND) If $a \in F(\theta)$ and $a_i(\theta)b \ \forall i \in I$, then $b \in F(\theta)$.

The first condition is a rather modest requirement on the richness of the domain of F . The second is a non-discrimination property which says that if everybody considers two allocations to be indifferent and one allocation belongs to the SCR, then it must be the other. Now we have the following:

Theorem 5 *Let F be a SCR satisfying L and ND and such that:*

1. *It is Nash implementable.*
2. *It selects individually rational allocations.*

Then, if x is a Walrasian allocation at θ , $x \in F(\theta)$.

Proof (Sketch; see Thomson 1985 for details) Take an economy θ . Let x be a Walrasian allocation for θ . Consider a new economy, called θ^L , where the marginal rates of substitution among goods are constant and equal to a vector of Walrasian prices. By individual rationality, F must select an allocation which is indifferent to x . By ND, $x \in F(\theta^L)$. Since F is Nash

implementable, it satisfies M. Now since $xR_i(\theta^L)b \rightarrow xR_i(\theta)b \ \forall i \in I$, by M, $x \in F(\theta)$.

Thus, under weak conditions, Walrasian allocation is always in the set of those selected by a monotonic SCR. And these allocations may fail to satisfy properties of fairness or justice as pointed out by the critics of the market. Under stronger assumptions, the converse is also true, i.e., only Walrasian allocations can be selected by a Nash-implementable SCR, Hurwicz (1979). Also, T. 5 has the following unpleasant implication:

Theorem 6 *There is no SCF in exchange economies such that:*

1. *It is Nash implementable.*
2. *It selects individually rational allocations.*
3. *ND holds.*
4. *It is defined on all exchange economies.*

Proof T. 5 implies that any Walrasian allocation belongs to the allocations selected by F . Since Walrasian equilibrium is not unique for some economies in the domain, hence the result.

T. 6 has a counterpart in social choice domains, Muller and Satterthwaite (1977).

Theorem 6' *There is no SCF in a social choice domain such that:*

1. *It is monotonic.*
2. *It is not dictatorial.*
3. *Its range is A with $\#A > 2$.*
4. *It is defined on all possible preferences.*

An implication of T. 6–6' is that single-valued SCR are still problematic. But the consideration of multivalued SCR brings a new problem: the existence of several Nash equilibria. For instance, if $a, b \in F(\theta)$ with a and b being efficient allocations, agents play a kind of “battle of the sexes” game with no clear results. Moreover, the Nash equilibrium in mixed strategies may yield allocations outside $F(\theta)$ (the concern about mixed strategy equilibria was first raised by Jackson 1992).

Now let us come to the good news. Firstly, the ND condition, which is essential for T. 5 to hold, is not as harmless as it appears to be. For instance, it

is not satisfied by the envy-free SCR; see Thomson (1987) for a discussion. Secondly, there are perfectly reasonable SCR which are monotonic: We have already encountered the constrained Walrasian SCR. Also any SCR selecting interior allocations in $\mathcal{LA}$ when preferences are von Neumann-Morgenstern is monotonic. In the domain of exchange economies with strictly increasing preferences, the core and the envy-free SCR are also monotonic. In domains where indifference curves only cross once – the single-crossing condition – monotonicity vacuously holds. So monotonicity, restrictive as it is, is worth a try. But before this, let us introduce a new assumption:

- (NVP) A SCR f satisfies no veto power if $\forall \theta \in \Theta$,

$$\{aR_i(\theta)b, \forall b \in A, \text{ for at least } n-1 \text{ agents}\} \\ \rightarrow a \in F(\theta)$$

In other words, if there is an allocation which is top-ranked by, at least, $n-1$ agents, NVP demands that this allocation belongs to the SCR. This sounds like a reasonable property for large n . Also in exchange economies with strictly increasing preferences and more than two agents, NVP is vacuously satisfied because there is no top allocation for $n-1$ agents.

The following positive result, a relief after so many negative results, was stated and proved by Maskin (1977), although his proof was incomplete:

Theorem 7 *If a SCR satisfies M and NVP is Nash implementable when $n > 2$.*

Proof (Sketch) Consider the following mechanism. $M_i - \Theta \times A \times \mathbb{N}$ where $\mathbb{N}$ is the set of natural numbers. The outcome function has three parts:

Rule 1 (unanimity). If m is such that all agents announce the same state of the world, θ , the same allocation a with $a \in F(\theta)$ and the same integer, then $g(m) = a$.

Rule 2 (one dissident). If there is only one agent whose message is different from the rest, this agent can choose any allocation that leaves her worse off, according to her preference as announced by others.

Rule 3 (any other case). $a \in g(m)$ if a was announced by the agent who announced the highest integer (ties are broken by an arbitrary rule).

Let us show that such a mechanism implements any SCR with the required conditions. Clearly, if the true state is $\tilde{\theta}$, $m_i = (\tilde{\theta}, a, 1)$ with $a \in F(\tilde{\theta})$ is a Nash equilibrium since no agent can gain by saying otherwise, so Condition 1 in the definition of Nash implementation holds. Let us now prove that Condition 2 there also holds. Suppose we have a Nash equilibrium in Rule 1. Could it be an “untruthful” equilibrium? If so, we have two cases. Either the announced preferences are a monotonic transformation of preferences at $\tilde{\theta}$, in which case, M implies that the announced allocation is also optimal at $\tilde{\theta}$. If they are not, there is an agent who can profitably deviate. Clearly, if equilibrium occurs in Rule 2, with, say, agent i as the dissident, any agent other than i can drive the mechanism to Rule 3, so it must be that all these agents are obtaining their most preferred allocation, which by NVP belongs to $F(\tilde{\theta})$. An equilibrium in Rule 3 implies that all agents are obtaining their most preferred allocation which, again by NVP, belongs to $F(\tilde{\theta})$.

The interpretation of the mechanism given in the proof of T. 7 is that if everybody agrees on the state and the allocation is what the planner wants, this allocation is selected. If there is a dissident (a term due to Danilov 1992), she can make her case by choosing an allocation (a “test allocation”) in her lower contour set, as announced by others. Finally, with more than one dissident, it is the jungle! Any agent can obtain her most preferred allocation by the choice of an integer. Typically, there is no equilibrium in this part of the mechanism. Notice that (M) is just used to eliminate unwanted equilibria.

The mechanism is an “augmented” revelation mechanism (a term due to Mookherjee and

Reichelstein 1990), where the announcement of the state is complemented with the announcement of an allocation – this can be avoided if the SCR is single valued – and an integer. The final proof of T. 7 was done independently by Williams (1986), Repullo (1987), Saijo (1988), and McKelvey (1989).

The case of two agents is more complicated because when an agent deviates from a common announcement and becomes a dissident, she converts the other agent into another dissident! As in T. 3, W does the job, i.e., any SCR satisfying M , NVP , and W is Nash implementable; see Moore and Repullo (1990) and Dutta and Sen (1991b) for a full characterization. Again, the cases of two agents and more than two agents are different. In some areas of mathematics, such as statistics and differential equations, the cases of two dimensions and more than two dimensions are also different. The relationship of these with the findings of implementation is not yet fully explored; see Saari (1987).

Under asymmetric information, M is substituted by a – rather ugly – *Bayesian monotonicity* (BM) condition which is a generalization of M to these environments. BM is again necessary and, in conjunction with some technical conditions plus incentive compatibility, sufficient for implementation in BE. The interested reader can do no better than to read the account of these matters in Palfrey (2002). It must be remarked that many well-known SCR – including Arrow-Debreu contingent commodities and some efficient SCR – do not satisfy BM and thus cannot be implemented in BE. However, the rational expectations equilibria and the (interim) envy-free SCR satisfy BM; see Palfrey and Srivastava (1987).

T. 7 was the first positive finding of implementation theory. And it prompted researchers to be more ambitious: Can we implement without monotonicity? An interesting observation, due to Matsushima (1988) and Abreu and Sen (1991), is that if agents have preferences representable by von Neumann-Morgenstern utility functions, any SCR can be “virtually implemented” in the sense that the set of allocations yielded by Nash equilibria is arbitrarily close to the set of desired allocations. This is because, as we saw before, any SCR mapping in the interior of $\mathcal{L}A$ is

monotonic. Thus, allocations in the boundary can be arbitrarily approximated by allocations in the interior.

A more satisfying approach was introduced by Moore and Repullo (1988) by introducing subgame perfection as the solution concept. It is not possible to explain fully this approach here because it would take us too far; in particular, the notion of a mechanism must be generalized to “stage mechanism.” Instead, we give a result that conveys the force of subgame perfect implementation. It refers to public good economies with quasi-linear utility functions – where under dominant strategies the set of economies with inefficient outcomes is large- and with two individuals – where Nash implementability is harder to obtain.

Suppose that utility functions read $U_i = V(y, \theta_i) + m_i$ where $y \in Y \subseteq \mathfrak{R}$, $\theta_i \in \Theta_i$, with $\#\Theta_i < \infty$ and $m_i \in \mathfrak{R}$, $i = 1, 2$. The set of allocations $\{(y, m_1, m_2) \in Y \times \mathfrak{R}^2 / m_1 + m_2 \leq \omega\}$ where ω are the endowments of “money.” Moore and Repullo (1988) proved the following:

Theorem 8 *Any SCF is implementable in subgame perfect equilibrium in the domain of economies explained above.*

Moore and Repullo proved that many SCR which could not be implemented in Nash equilibrium can be implemented in subgame perfect equilibrium. This is because subgames can be designed to kill unwanted equilibria without using monotonicity. Their result was improved upon by Abreu and Sen (1991). The problem with this approach is that the concept of subgame perfection is problematic because it requires that, no matter what has happened in past, in the remaining subgame, players are rational, even if this subgame was attained because some players made irrational choices.

The Moore-Repullo result was not only important by itself but it opened the way to the consideration of other equilibrium concepts that allow very permissive results. For instance, Palfrey and Srivastava (1991) proved the following result:

Theorem 8' *Any SCR satisfying NVP is implementable in undominated Nash equilibrium.*

At this point, it seemed that by invoking the adequate refinement of Nash equilibrium, any SCR could be implemented. But the implementing mechanisms were getting weird and some people were beginning to get suspicious. Why and how is discussed in the next section.

Summing up the results obtained here, we have the following:

1. (Maskin) Monotonicity is a necessary and, in many cases, sufficient condition for implementation in Nash equilibrium; see T. 4 and 7. Similar results are obtained with Bayesian monotonicity in Bayesian equilibrium.
2. The monotonicity requirements are not harmless. Many solution concepts do not satisfy it. Even worse, monotonicity has some unpalatable consequences; see T. 5–6.
3. Monotonicity can be avoided by considering stage games or refinements of Nash equilibrium. Practically, any reasonable SCR can be implemented in this way; see T. 8–8'.

The Limits of Design

So far, we have assumed that there are no limits to what the designer can do. She can pick up any mechanism with no restrictions on its shape. This procedure, indeed, pushes the possibilities of design to the limit. But by doing this, we have learnt a good deal about the limitations of the theory of implementation. It is fair to say that today the consensus is that there are *some extra properties* which should be considered when designing an implementing mechanism. We review here five approaches to this question:

1. **Game-Theoretical Concerns.** Jackson (1992) was the first to point out that some mechanisms had unusual features from the point of view of game theory: Some subgames have no Nash equilibrium. Message spaces, which in the corresponding game become strategy spaces, are unbounded or open. Thus, in the integer game considered in T. 7, if agents eliminate dominated strategies, each integer is dominated by the next highest one, and no integer is undominated: Those agents who eliminate dominated strategies are unable to make a
- choice. These constructions eliminate unwanted equilibria, which, as we saw before, is the problem with Nash implementation. Jackson illustrates his point by showing that under no restrictions on mechanisms, any SCR can be implemented in undominated strategies, a weak solution concept. Then, he requires that the mechanism be *bounded* in the following sense: Whenever a strategy m_i is dominated, there is another strategy dominating m_i and which is undominated. He shows that implementation in undominated strategies with bounded mechanisms results many times in incentive compatibility, which as we saw in section “[The Main Concepts](#)” is a hard requirement. This shows the bite of the boundedness assumption. However, in the case of implementation with undominated Nash equilibrium, the boundedness assumption has little impact; see Jackson et al. (1994) and Sjöström (1994). The first of these papers introduced a related requirement, the *best response property*: For every strategy played by the other agents, each agent has a best response.
2. **Natural Mechanisms.** Given that we have run so far from the kind of mechanism we are used to, it seems reasonable to ask what can be implemented by mechanisms that resemble real-life mechanisms. These mechanisms must be simple too because simplicity is an important characteristic in practice. Let us call them *natural mechanisms*. Dutta et al. (1995) consider mechanisms in which messages are prices and quantities and thus resemble market mechanisms. Their approach was refined by Saijo et al. (1996) who demanded the best response property as well. They showed that several well-known SCR, such as the (constrained) Walrasian, are implementable in Nash equilibrium. Beviá et al. (2003) showed that in Bertrand-like market games, the Walrasian SCR is implementable in Nash and strong equilibrium, showing that the fear of coalitions destabilizing market outcomes is, at least, partially unwarranted. Sjöström (1996) considered quantity mechanisms, reminiscent of those used by Soviet planners, with negative results about what these mechanisms can

achieve. In public good economies, Corchón and Wilkie (1996) and Peleg (1996) introduced a market mechanism implementing Lindahl allocations in Nash and strong equilibrium. The mechanism works because Lindahl prices have to add up to the marginal cost. If an agent pretends to free ride, she decreases the quantity of the public good. Here, contrary to Samuelson's dictum, it is in the selfish interest of each person to give true signals. Pérez-Castrillo and Wettstein (2002) offered a bidding mechanism that implements efficient allocations when choosing between a finite numbers of public projects. They also applied these ideas to extend the Shapley value to more general environments, Macho-Stadler et al. (2007).

3. **Credibility.** Another implicit assumption is that once the mechanism is in place, there is no way to stop it. Thus, if for some m , $g(m)$ is a "universally worst outcome," the planner has to deliver this allocation even if she is trying to implement a Pareto efficient allocation. Is this a credible procedure? In many cases, if the planner is a real person, it seems that she would do her best to avoid $g(m)$! Here we have two possibilities: either we identify additional constraints on the planner that look reasonable or we jump to model the planner as a full-fledged player. The first road leads us to identify a subset of allocations of A , say X , which can never be used by the mechanism. For instance, in Chakravorty et al. (2006) X is the set of allocations that are never selected by the SCR for some state of the world, i.e., $X = \{a \in A / \nexists \theta \in \Theta, a \in F(\theta)\}$. The motivation for this definition is that it hardly seems credible that the planner can choose an allocation that is never intended to be implemented. Redefining the allocation set as $A \equiv A \setminus X$, the definitions of a mechanism and an implementable SCR can be easily translated in this framework. However, depending on the domain, SCR that are monotonic when defined on A are no longer monotonic when defined on A , for instance, the (constrained) Walrasian SCR. Thus, these SCR cannot be implemented when the planner can only use allocations in A . A weakness of

this approach is that the list of reasonable constraints on allocations may be large. The second possibility drives us to model implementation as a signaling game where the planner receives signals – messages – from the agents, updates her beliefs, and then chooses an allocation which maximizes her expected utility (Baliga et al. 1997). Again, some SCR that are Nash implementable are not implementable in this framework. However, in this case, there are SCR that are not Nash implementable but are implementable in this framework. This is because the model takes a basic assumption of game theory to the limit, namely, that agents know the strategies of other players. In this case, the planner knows if a report on agents' types is truthful or not before the allocation is delivered.

4. **Renegotiation.** Another strong assumption is that the mechanism prescribes actions that cannot be changed by agents. This contradicts experiences such as black markets where agents trade on the existing goods (Hammond 1987). A way of modeling this is to assume that agents are able to renegotiate some allocations (Maskin and Moore 1999). Renegotiation in a different context was considered by Rubinstein and Wolinsky (1992). Assuming that agents have complete information, this is formalized by means of the concept of a reversion function. This function, say r , maps each allocation and each state of the world into a new allocation, i.e., $r : A \times \Theta \rightarrow A$. The reversion function induces new preferences, called *reverted preferences* (this is the "translation principle" in Maskin and Moore (1999)). Notice that reverted preferences are state-dependent even if preferences are not. Formally, given a reversion function r , the *reversion* of $R(\theta)$, denoted by $R^r(\theta)$, is defined as $aR_i^r(\theta)b \Leftrightarrow r(a, \theta)R_i(\theta)r(b, \theta), \forall a, b \in A, \forall i \in I$. Given a reversion function r , we can interpret that agents' preferences are the reverted preferences. Then, all definitions given before can be adapted to this case. Again, SCR that were monotonic there are not so in this framework and vice versa. See Jackson and Palfrey (2001) for applications. An extension to the case where there are several renegotiation functions is given by Amorós

(2004). A weakness of this approach is that it models renegotiation as a “black box.”

5. **Multiple Implementation.** Maskin (1985) was the first to realize that the notion of implementation requires the planner to know the solution concept used by the agents to analyze the game. He proposed the notion of “double implementation” where a SCR was implemented at the same time in Nash and strong equilibria. He showed that many Nash-implementable SCR indeed are doubly implementable. We have seen in Point 2 above that the (constrained) Walrasian and Lindahl SCR are doubly implementable by natural mechanisms. They are also doubly implementable by abstract mechanisms, Schmeidler (1980). Double implementation also occurs with several solutions to the problem of the commons, Shin and Suh (1997), and Pigouvian taxes, Alcalde et al. (1999). Yamato (1993) introduced another type of double implementation by requiring implementation in Nash and undominated Nash equilibria (1993). He showed that in a large class of exchange economies with at least three agents, monotonicity is necessary and sufficient for double implementation. Saijo et al. (2007) considered implementation in dominant strategies and Nash equilibrium. Clearly, other variations of the idea of double implementation are possible; see Point 4 in section “[Unsolved Issues and Further Research](#)” below.

Summing up, it is now clear that implementing mechanisms cannot be just “anything.” Their features matter. Demanding that mechanisms satisfy the best response property, be simple, do not use extreme allocations, and be robust to the possibility of renegotiation and implement in several equilibrium concepts makes our lives more difficult but makes our models a great deal better.

Unsolved Issues and Further Research

1. **Implementation with State-Dependent Feasible Sets.** A motivation of implementation theory was to study the possibility of socialism.

However, *all* the results presented in this survey refer to environments where the feasible set is given, a far cry from any kind of planning procedure. In fact, there are only a handful of papers dealing with implementation when the feasible set is unknown: Postlewaite (1979) and Sertel and Sanver (1999) studied manipulation of endowments. Hurwicz et al. (1995) studied implementation assuming that endowments/production possibilities can be hidden or destroyed but never exaggerated. Instead of a mechanism, we have a collection of state-dependent mechanisms each meant for an economy. After the mechanism is played, production capabilities are shown, e.g., endowments are put on the table. This idea was worked out in a series of papers by Hong on private good economies, Hong (1998), and by Tian on public good economies, Tian and Li (1995). Serrano and Vohra (1997) worked out implementation of the core and Dagan et al. (1999) of taxation methods. And that is all folks! Why has such an important issue been almost neglected? My explanation is that the proposed mechanisms are difficult to understand. Another approach has been tried by Corchón and Triossi (2011) where a reversion function takes care of restoring feasibility when messages lead to unfeasible allocations. The approach is tractable and simpler but relies on the black box of the renegotiation function.

2. **Sociological Factors/Bounded Rationality.** So far, all the solution concepts describing the behavior of agents are game theoretical. In recent years, we have seen a host of equilibrium concepts based on “irrational” agents. It would be interesting to see what SCR can be implemented with these forms of behavior. Eliaz (2002) considers “fault-tolerant” implementation where a subset of players (“faulty players”) fail to achieve their optimal strategies. Under complete information, no veto power and a strong form of monotonicity are sufficient for implementation when the number of faulty players is less than $n/2 - 1$, $n > 2$. Matsushima (2008) and Dutta and Sen (2012) show that a small preference for honesty is sufficient to knock down unwanted equilibria.

Corchon and Herrero (2004) show that “decent behavior” can also be used to dispense with monotonicity.

3. **Dynamic Implementation.** The theory presented here is static but there are some papers dealing with implementation in dynamic setups. We mention a few: Freixas et al. (1985) studied the “ratchet effect”, where firms underproduce for fear of being asked to do too much in the future. Kalai and Ledyard (1988) showed that if the planner is sufficiently patient, every SCR is dominant strategy implementable. Burguet (1990/1994) showed that the revelation principle does not hold when outcomes are chosen in several periods. Candel (2004) proved a revelation principle in a model where a public good is produced in two periods. Cabrales (1999) and Sandholm (2007) studied implementation in an evolutionary setting. A related topic is that of complexity; see Conitzer and Sandholm (2002). Lee and Sabourian (2011) studied implementation with infinitely lived agents whose preferences are determined randomly in each period. A SCR is repeated-implementable in Nash equilibrium if there exist a sequence of (possibly history-dependent) mechanisms such that (i) its equilibrium set is nonempty and (ii) every equilibrium outcome corresponds to the desired social choice at every possible history of past play and realizations of uncertainty. They show that, essentially, a SCR is repeated-implementable if and only if it is efficient.
4. **Robustness Under Incomplete Information.** When designing a mechanism, sometimes the planner does not know the structure of information. In this case, a mechanism must be implemented regardless of the structure of information, i.e., priors of agents, type spaces, etc. Corchón and Ortuño-Ortín (1995) approached the problem by assuming that the economy is composed of “islands” and that there is complete information inside each island. A mechanism robustly implements a SCR if it does it in BE for every possible prior (compatible with the island assumption) and in uniform Nash equilibrium. The latter requires that an equilibrium strategy for an agent must be the best reply to what other

agents in the island play and to any possible message sent by agents outside the island when they follow their equilibrium strategies (D’Aspremont and Gerard-Varet 1979). They showed that any SCR satisfying M and NVP is robustly implementable (a later contribution by Yamato 1994 showed that robust and Nash implementation coincide in this framework). The same concern has been approached in a series of papers by Bergemann and Morris (see, e.g., 2005) where they ask SCR to be implemented whatever the players’ beliefs and higher-order beliefs about other player’s types. Artemov et al. (2007) require implementation for the payoff type space and the space of first-order beliefs about other agents’ payoff types. They obtain very permissive results.

In a different vein, Koray (2005) has argued that, since priors are not contractible, the regulator needs to be regulated in order to stop her from manipulating the priors. He shows that the outcomes of this game vary over a wide spectrum. Again the need of prior-free implementation is clear.

Answers to the Questions

1. Yes. We already saw in 4.3, Point 2, that “Bertrand-like” mechanisms implement the constrained Walrasian SCR in Nash and strong equilibrium. But this is not all: Schmeidler (1980) exploited the connection between price-taking, which underlies Walrasian equilibrium, and “strategy-taking”, which underlies Nash and strong equilibrium and obtained double implementation by a mechanism which does not resemble the market. Implementation of the Lindahl SCR by an abstract mechanism was obtained by Walker (1981) building on previous papers by Groves and Ledyard and Hurwicz. Unfortunately, these positive results turn negative when we consider Arrow-Debreu contingent commodities, Chattopadhyay et al. (2000) and Serrano and Vohra (2001).
2. A merger affects social welfare in two ways: positively, from cost savings, and negatively, from restricting competition. The first effect is

uncertain, and by now, I do not have to convince you that we should take with utmost caution all announcements made by firms concerning cost savings. Corchón and Faulí-Oller (2004) show that under a condition that is fulfilled in several standard IO models, the SCR that maximize social surplus can be implemented by a dominance-solvable mechanism with budget balance.

3. There is a very simple mechanism which attains maximum surplus, Loeb and Magath (1979). But in this mechanism, the monopolist receives all the surplus and the demand function must be known by the planner. These points were worked out by subsequent contributions from Baron and Myerson, Lewis and Sappington, Sibley, and others.
4. By now, the reader should know the difficulties of implementing efficient public decisions. When information is exclusive, this is impossible, even though an approximate efficient decision can be obtained when the number of agents is large. When information is complete, we have seen several examples of mechanisms implementing efficient outcomes.
5. There is no difference between implementing market and fair outcomes. Both have to pass the same tests, i.e., incentive compatibility, monotonicity, and simplicity/credibility of design. In exchange economies, Thomson (2005) presents a simple and elegant mechanism that implements envy-free allocations in Nash equilibrium. In cooperative production, Corchón and Puy (2002) presented a family of mechanisms that implement in Nash equilibrium any efficient SCR where the distribution of rewards is a continuous function of efforts.
6. Yes! An uninformed planner can set up a mechanism that yields efficient outcomes in circumstances where the market yields inefficient allocations, i.e., under externalities or public goods; see Point 5 in section “[The Limits of Design](#)” above. All we need is nonexclusive information and that the SCR be monotonic; the latter requirement can be skipped under refinements of Nash equilibrium.
7. Not completely. Suppose complete information among three or more judges and that they

all perceive the same quality of a given performance. Clearly, truth is an equilibrium, because if all judges minus you tell the truth, you cannot change the outcome by saying something different. Unfortunately, any unanimous announcement is also an equilibrium by the same reason. Thus, we are in a situation akin to T. 3. Fortunately, if preferences of judges fulfill certain restrictions, full implementation of the true ranking of ice skaters is possible, because monotonicity and no veto power hold so T. 7 applies, Amorós et al. (2002). Amorós (2009) found necessary and sufficient conditions on the preferences of the judges for the socially optimal ranking to be the Nash implementable. Later on, Amorós (2016) presented a natural mechanism to pick up only the deserving winner (not a complete ranking of the participants). If judges have differential information, the truth is no longer implementable as suggested by T. 2”. See Gerardi et al. (2005) for further insights and references on this problem.

8. ????? Do you think that we have all answers? This is just economics!!

Finally, I will tell you why I like implementation theory so much. Firstly, the implementation model solves the problems of the general equilibrium model mentioned in section “[Brief History of Implementation Theory](#),” namely, (1) it models a general economic system, (2) all variables are endogenously determined by the interaction of agents, and (3) agents’ incentives are carefully modeled and are taken fully into account. Secondly, the theory is not based on assumptions like convexity or continuity/differentiability which, no matter how much we are used to them, are very stringent. By the way, a beautiful paper by Lafont and Maskin (referenced in their 1982 survey) developed incentive compatibility in a differentiable framework.

Acknowledgments I am grateful to Pablo Amorós, Claude d’Aspremont, Carmen Beviá, Luis Cabral, Eric Maskin, Bernardo Moreno, Carlos Pimienta, Socorro Puy, Tömas Sjöström, William Thomson, Matteo Triossi, Galina Zudenkova, and an anonymous referee for helpful suggestions and to the Spanish Ministry of Education for

financial support under grant SEJ2005-06167. I also thank the Department of Economics, Stern School of Business, NYU, for their hospitality while writing the first draft of this survey. This survey was dedicated to Leo Hurwicz to celebrate his Nobel Prize and to the memory of those who contributed to the area and are no longer with us: Louis-André Gerard-Varet, Jean-Jacques Laffont, Richard McKelvey, and Murat Sertel.

Bibliography

- Abreu D, Sen A (1990) Virtual implementation in Nash equilibrium. *Econometrica* 59:997–1021
- Abreu D, Sen A (1991) Subgame perfect implementation: a necessary and almost sufficient condition. *J Econ Theory* 50:285–299
- Akerlof G The market for lemons: qualitative uncertainty and the market mechanism. *Q J Econ* 84:488–500
- Alcalde J, Corchón L, Moreno B (1999) Pigouvian taxes: a strategic approach. *J Public Econ Theory* 1(2):271–281
- Amoros P (2004) Nash implementation and uncertain renegotiation. *Game Econ Behav* 49:424–434
- Amorós P (2009) Eliciting socially optimal rankings from unfair jurors. *J Econ Theory* 144:1211–1226
- Amorós P (2016) Subgame perfect implementation of the deserving winner of a competition with natural mechanisms. *Math Soc Sci* 83:44–57
- Amoros P, Corchón L, Moreno B (2002) The scholarship assignment problem. *Game Econ Behav* 38:1–18
- Arrow K (1977) The property rights doctrine and demand revelation under incomplete information. Technical report No. 243, IMSSS. Stanford University, Aug 1977
- Artemov G, Kunimoto T, Serrano R (2007) Robust virtual implementation with incomplete information: towards a reinterpretation of the Wilson doctrine. W.P. 2007-06. Brown University
- Baliga S, Sjöström T (2007) Mechanism design: recent developments. In: Blume L, Durlauf S (eds) *The new palgrave dictionary of economics*, 2nd edn
- Baliga S, Corchón L, Sjöström T (1997) The theory of implementation when the planner is a player. *J Econ Theory* 77:15–33
- Barberá S (1983) Strategy-proofness and pivotal voters: a direct proof of the Gibbard-Satterthwaite theorem. *Int Econ Rev* 24:413–417
- Barberá S, Peleg B (1990) Strategy-proof voting schemes with continuous preferences. *Soc Choice Welf* 7:31–38
- Barberá S, Sonnenschein H, Zhou L (1991) Voting by committees. *Econometrica* 59(3):595–609
- Barberá S, Berga D, Moreno B (2010) Individual versus group strategy-proofness: when do they coincide? *J Econ Theory* 145:1648–1674
- Barberá S, Berga D, Moreno B (2016) Group strategy-proofness in private good economies. *Am Econ Rev* 106(4):1073–1099
- Barone E (1908) The ministry of production in a collectivist state. Translated from the Italian and reprinted in F. A. von Hayek *Collectivist Economic Planning*, Routledge and Keegan, 1935
- Benoit J-P (2000) The Gibbard-Satterthwaite theorem: a simple proof. *Econ Lett* 69:319–322
- Bergemann D, Morris S (2005) Robust mechanism design. *Econometrica* 73:1771–1813
- Beviá C, Corchón L (1995) On the generic impossibility of truthful behavior. *Economic Theory* 6:365–371
- Beviá C, Corchón L, Wilkie (2003) Implementation of the Walrasian correspondence by market games. *Rev Econ Des* 7:429–442
- Burguet R (1990) Revelation in informational dynamic settings. *Econ Lett* 33:237–239. Corrigendum (1994), 44:451–452
- Cabrales A (1999) Adaptive dynamics and the implementation problem with complete information. *J Econ Theory* 86:159–184
- Candel F (2004) Dynamic provision of public goods. *Economic Theory* 23:621–641
- Chakravorty B, Corchón L, Wilkie S (2006) Credible implementation. *Game Econ Behav* 57:18–36
- Chattopadhyay S, Corchón L, Naeve J (2000) Contingent commodities and implementation. *Econ Lett* 68:293–298
- Clarke E (1971) Multipart pricing of public goods. *Public Choice* 19–33
- Conitzer V, Sandholm T (2002) Complexity of mechanism design. In: *Proceedings of the 18th annual conference on uncertainty in artificial intelligence (UAI-02)*, Edmonton
- Corchón L (1996) *The theory of implementation of socially optimal decisions in economics*. St. Martin's Press, New York
- Corchón L, Faulí-Oller R (2004) To merge or not to merge: that is the question. *Rev Econ Des* 9:11–30
- Corchón L, Herrero C (2004) A decent proposal. *Span Econ Rev* 6(2):107–125
- Corchón L, Ortuño-Ortín I (1995) Robust implementation under alternative information structures. *Econ Des* 1:159–171
- Corchón L, Puy S (2002) Existence and Nash implementation of efficient sharing rules for a commonly owned technology. *Soc Choice Welf* 19:369–379
- Corchón L, Triossi M (2011) Implementation with renegotiation when preferences and feasible sets are state dependent. *Soc Choice Welf* 36(2):179–198
- Corchón L, Wilkie S (1996) Doubly implementing the ratio correspondence by a market mechanism. *Rev Econ Des* 2:325–337
- d'Aspremont C, Gerard-Varet L-A (1975) Individual incentives and collective efficiency for an externality game with incomplete information. CORE DP 7519
- d'Aspremont C, Gerard-Varet L-A (1979) Incentives and incomplete information. *J Public Econ* 11:25–45
- Dagan N, Serrano R, Volij O (1999) Feasible implementation of taxation methods. *Rev Econ Des* 4:57–72
- Danilov V (1992) Implementation via Nash equilibria. *Econometrica* 60(1):43–56
- Dasgupta P, Hammond P, Maskin E (1979) The implementation of social choice rules: some results on incentive compatibility. *Rev Econ Stud* 46:185–216
- Dutta B, Sen A (1991a) Implementation under strong equilibrium. A complete characterization. *J Math Econ* 20:49–67

- Dutta B, Sen A (1991b) A necessary and sufficient condition for two-person Nash implementation. *Rev Econ Stud* 58:121–128
- Dutta B, Sen A (2012) Nash implementation with partially honest individuals. *Game Econ Behav* 74(1):154–169
- Dutta B, Sen A, Vohra R (1995) Nash implementation through elementary mechanisms in economic environments. *Econ Des* 1:173–204
- Eliasz K (2002) Fault tolerant implementation. *Rev Econ Stud* 69:589–610
- Freixas X, Guesnerie R, Tirole J (1985) The ratchet effect. *Rev Econ Stud* 52:173–191
- Gerardi D, McLean R, Postlewaite A (2005) Aggregation of expert opinions. Cowles Foundation Discussion Paper # 1503
- Gibbard A (1973) Manipulation of voting schemes: a general result. *Econometrica* 41:587–602
- Green J, Laffont J-J (1979) Incentives in public decision making. North Holland, Amsterdam
- Gresik T, Satterthwaite M (1989) The rate at which a simple market converges to efficiency as the number of traders increases. *J Econ Theory* 48:304–332
- Groves T (1973) Incentives in teams. *Econometrica* 41:617–631
- Hammond P (1987) Markets as constraints: multilateral incentive compatibility in continuum economies. *Rev Econ Stud* 54:399–412
- Harris M, Townsend R (1981) Resource allocation under asymmetric information. *Econometrica* 49:33–64
- Harsanyi J (1967/1968) Games with incomplete information played by ‘Bayesian’ players. Parts I, II and III. *Manag Sci* 14:159–182, 320–334 and 486–502
- Hong L (1998) Feasible Bayesian implementation with state dependent feasible sets. *J Econ Theory* 80:201–221
- Hurwicz L (1959) Optimality and informational efficiency in resource allocation processes. In: Arrow KJ (ed) *Mathematical methods in the social sciences*. Stanford University Press, Stanford, pp 27–46
- Hurwicz L (1972) On informationally decentralized systems. In: Radner R, McGuire CB (eds) *Decision and organization: a volume in honor of Jacob Marshak*. North-Holland, Amsterdam, pp 297–336
- Hurwicz L (1979) On allocations attainable through Nash equilibria. *J Econ Theory* 21:40–65
- Hurwicz L, Walker M (1990) On the generic non-optimality of dominant strategy mechanisms. *Econometrica* 58(3):683–704
- Hurwicz L, Maskin E, Postlewaite A (1995) Feasible Nash implementation of social choice rules when the designer does not know endowments or production set. In: Ledyard J (ed) *The economics of informational decentralization: complexity, efficiency and stability*. Kluwer, Dordrecht
- Jackson M (1991) Bayesian implementation. *Econometrica* 59:461–477
- Jackson M (1992) Implementation in undominated strategies: a look to bounded mechanisms. *Rev Econ Stud* 59:757–775
- Jackson MO (2001) A crash course in implementation theory. *Soc Choice Welf* 18:655–708
- Jackson MO, Palfrey T (2001) Voluntary implementation. *J Econ Theory* 98:1–25
- Jackson MO, Palfrey T, Srivastava S (1994) Undominated Nash implementation in bounded mechanisms. *Game Econ Behav* 6:474–501
- Kalai E, Ledyard J (1988) Repeated implementation. *J Econ Theory* 83:308–317
- Klemperer P (1999) Auction theory, a guide to the literature. *J Econ Surv* 13:227–268
- Koray S (2005) The need of regulating a Bayesian regulator. *J Regul Econ* 28:5–21
- Krishna V (2002) Auction theory. Academic, San Diego
- Krishna V, Perry M (1997) Efficient mechanism design. Unpublished paper, Penn State University
- Laffont J-J, Martimort D (2001) *The theory of incentives: the principal-agent model*. Princeton University Press, Princeton
- Laffont J-J, Maskin E (1982) The theory of incentives: an overview. In: Hildenbrand W (ed) *Advances in economic theory*, 4th World Congress of the Econometric Society. Cambridge University Press, New York
- Lange O (1936/1937) On the economic theory of socialism. *Rev Econ Stud* 4:53–71, 123–142
- Ledyard J, Roberts J (1974) On the incentive problem with public goods. Discussion Paper 116, Centre for Mathematical Studies in Economics and Management Science. Northwestern University
- Lee J, Sabourian H (2011) Efficient repeated implementation. *Econometrica* 79(6):1967–1994
- Loeb M, Magath W (1979) A decentralized method for utility regulation. *J Law Econ* 22:399–404
- Ma A, Moore J, Turnbull S (1988) Stop agents from cheating. *J Econ Theory* 46:355–372
- Macho-Stadler I, Pérez-Castrillo D, Wettstein D (2007) Sharing the surplus: an extension of the Shapley value for environments with externalities. *J Econ Theory* 135(1):339–356
- Makowski L, Mezzetti C (1993) The possibility of efficient mechanisms for trading an indivisible object. *J Econ Theory* 59:451–465
- Maskin E (1985) The theory of implementation in Nash equilibrium: a survey. In: Hurwicz L, Schmeidler D, Sonnenschein H (eds) *Social goals and social organization*. Cambridge University Press, Cambridge, UK, pp 173–204
- Maskin E (1999) Nash equilibrium and welfare optimality. *Rev Econ Stud* 66(1):23–38. Circulating in working paper version since 1977
- Maskin E, Moore J (1999) Implementation with renegotiation. *Rev Econ Stud* 66:39–56
- Maskin E, Sjöström T (2002) Implementation theory, Chapter 5. In: Arrow KJ, Sen AK (eds) *Handbook of social choice and welfare*. Elsevier
- Matsushima H (1988) A new approach to the implementation problem. *J Econ Theory* 45:128–144
- Matsushima H (2008) Role of honesty in full implementation. *J Econ Theory* 139:353–359
- McKelvey R (1989) Game forms for Nash implementation of general social choice correspondences. *Soc Choice Welf* 6:139–156
- Mirrless J (1971) An exploration in the theory of optimum income taxation. *Rev Econ Stud* 38:175–208

- Mookherjee D, Reichelstein S (1990) Implementation via augmented revelation mechanisms. *Rev Econ Stud* 57:453–475
- Moore J (1992). Implementation, contracts and renegotiation in environments with complete information. In: Laffont J-J (ed) *Advances in economic theory*, 4th World Congress of the Econometric Society, vol I. Cambridge University Press, Cambridge, UK, pp 182–282
- Moore J, Repullo R (1988) Subgame perfect implementation. *Econometrica* 56:1191–1220
- Moore J, Repullo R (1990) Nash implementation: a full characterization. *Econometrica* 58:1083–1089
- Moulin H (1979) Dominance solvable voting schemes. *Econometrica* 47(6):1337–1351
- Muller E, Satterthwaite M (1977) The equivalence of strong positive association and strategy-proofness. *J Econ Theory* 14:412–418
- Myerson R (1979) Incentive compatibility and the bargaining problem. *Econometrica* 47:61–73
- Myerson R (1981) Optimal auction design. *Math Oper Res* 6:58–73
- Myerson R (1985) Bayesian equilibrium and incentive compatibility: an introduction, Chapter 8. In: Hurwicz L, Schmeidler D, Sonnenschein H (eds) *Social goals and social organization*. Cambridge University Press, Cambridge, UK
- Myerson R, Satterthwaite MA (1983) Efficient mechanisms for bilateral trading. *J Econ Theory* 29:265–281
- Palfrey TR (2002) Implementation theory. In: Aumann RJ, Hart S (eds) *Handbook of game theory with economic applications*, vol III. Elsevier Science, New York, pp 2271–2326
- Palfrey T, Srivastava S (1987) On Bayesian implementable allocations. *Rev Econ Stud* 54:193–208
- Palfrey T, Srivastava S (1989) Implementation with incomplete information in exchange economies. *Econometrica* 57:115–134
- Palfrey T, Srivastava S (1991) Nash implementation using undominated strategies. *Econometrica* 59:479–501
- Peleg B (1996) Double implementation of the Lindahl equilibrium by a continuous mechanism. *Rev Econ Des* 2(1):311–324
- Pérez-Castrillo D, Wettstein D (2002) Choosing wisely: a multi-bidding approach. *Am Econ Rev* 92:1577–1587
- Postlewaite A. Manipulation via endowments. *Rev Econ Stud* 46:255–262
- Postlewaite A, Schmeidler D (1986) Implementation in differential information economies. *J Econ Theory* 39:14–33
- Prize Committee of the Royal Swedish Academy of Sciences (2007) Mechanism design theory. Royal Swedish Academy of Sciences, Stockholm
- Reiter S (1977) Information and performance in the (new) welfare economics. *Am Econ Rev* 67:226–234
- Repullo R (1986) On the revelation principle under complete and incomplete information. In: Binmore K, Dasgupta P (eds) *Economics organizations as games*. Basil Blackwell, Oxford
- Repullo R (1987) A simple proof of Maskin theorem on Nash implementation. *Soc Choice Welf* 4:39–41
- Roth AE (2008) What have we learned from market design? *Econ J* 118:285–310
- Rothchild M, Stiglitz J (1976) Equilibrium in competitive insurance markets: an essay on the economics of imperfect information. *Q J Econ* 90:629–650
- Rubinstein A, Wolinsky A (1992) Renegotiation-proof implementation and time preferences. *Am Econ Rev* 82:600–614
- Saari D (1987) The source of some paradoxes from social choice and probability. *J Econ Theory* 41:1–22
- Saijo T (1988) Strategy space reduction in Maskin's theorem. *Econometrica* 56:693–700
- Saijo T (1991) Incentive compatibility and individual rationality in public good economies. *J Econ Theory* 55:103–112
- Saijo T, Tatamitani Y, Yamato T (1996) Toward natural implementation. *Int Econ Rev* 37(4):949–980
- Saijo T, Sjöström T, Yamato T (2007) Secure implementation. *Theor Econ* 2:203–229
- Samuelson PA (1954) The pure theory of public expenditure. *Rev Econ Stat* 36:387–389
- Sandholm W (2007) Pigouvian pricing and stochastic evolutionary implementation. *J Econ Theory* 132:367–382
- Satterthwaite (1975) Strategy-proofness and arrows conditions: existence and correspondence theorems for voting procedures and social choice functions. *J Econ Theory* 10:187–217
- Schmeidler D (1980) Walrasian analysis via strategic outcome functions. *Econometrica* 48:1585–1593
- Sen A (2001) Another direct proof of the Gibbard-Satterthwaite theorem. *Econ Lett* 70:381–385
- Serizawa S (2002) Inefficiency of strategy-proof rules for pure exchange economies. *J Econ Theory* 106:219–241
- Serizawa S, Weymark J (2003) Efficient strategy-proof exchange and minimum consumption guarantees. *J Econ Theory* 109:246–263
- Serrano R (2004) The theory of implementation of social choice rules. *SIAM Rev* 46:377–414
- Serrano R, Vohra R (1997) Non cooperative implementation of the core. *Soc Choice Welf* 14:513–525
- Serrano R, Vohra R (2001) Some limitation of Bayesian virtual implementation. *Econometrica* 69:785–792
- Sertel M, Samver R (1999) Equilibrium outcomes of Lindahl-endowment pretension games. *Eur J Polit Econ* 15:149–162
- Shin S, Suh S-C (1997) Double implementation by a simple game form in the commons problem. *J Econ Theory* 77:205213
- Sjöström T (1993) Implementation in perfect equilibria. *Soc Choice Welf* 10:97–106
- Sjöström T (1994) Implementation in undominated Nash equilibrium without integer games. *Game Econ Behav* 6:502–511
- Sjöström T (1996) Implementation by demand mechanisms. *Econ Des* 1:343–354
- Spence M (1973) Job market signalling. *Q J Econ* 87:355–374

- Taylor FM (1929) The guidance of production in a socialist state. *Am Econ Rev* 19:1–8
- Thomson W (1985) Manipulation and implementation in economics. Unpublished manuscript, Rochester
- Thomson W (1987) The vulnerability to manipulative behavior of economic mechanisms designed to select equitable and efficient outcomes, Chapter 14. In: Groves T, Radner R, Reiter S (eds) *Information, incentives and economic mechanisms*. University of Minnesota Press, pp 375–396
- Thomson W (1996) Concepts of implementation. *Jpn Econ Rev* 47:133–143
- Thomson W (2005) Divide and permute. *Game Econ Behav* 52:186–200
- Tian G (1996) On the existence of optimal truth-dominant mechanisms. *Econ Lett* 53:17–24
- Tian G, Li Q (1995) On Nash-implementation in the presence of withholding. *Game Econ Behav* 9:222–233
- Vickrey W (1961) Counterspeculation, auctions and competitive sealed tenders. *J Financ* 16:8–37
- Walker M (1981) A simple incentive compatible scheme for attaining Lindahl allocations. *Econometrica* 49:65–73
- Williams S (1986) Realization of Nash implementation: two aspects of mechanism design. *Econometrica* 54:139–151
- Williams S (1999) A characterization of efficient, Bayesian incentive compatible mechanisms. *Economic Theory* 14:155–180
- Yamato T (1993) Double implementation in Nash and undominated Nash equilibria. *J Econ Theory* 59(2):311–323
- Yamato T (1994) Equivalence of Nash implementability and robust implementability with incomplete information. *Soc Choice Welf* 11:289–303

Two-Sided Matching Models

Marilda Sotomayor^{1,2} and Ömer Özak³

¹Department of Economics, University of Sao Paulo, Sao Paulo, Brazil

²EPGE Brazilian School of Economics and Finance, Sao Paulo, Brazil

³Department of Economics, Brown University, Providence, USA

Article Outline

Glossary

Definition of the Subject

Basic Definitions

A Brief Historical Account

Gale-Shapley Algorithm with the Colleges

Proposing to the Applicants

Gale-Shapley Algorithm with the Applicants

Proposing to the Colleges

Introduction

Discrete Two-Sided Matching Models

Continuous Two-Sided Matching Model With

Additively Separable Utility Functions

Hybrid One-to-One Matching Model

Incentives

Future Directions

Bibliography

Glossary

Achievable mate for agent y in a discrete two-sided matching model is any y 's of partner under some stable matching.

Additively separable preferences in a discrete two-sided matching model with sides F and W . Agent $f \in F$ has additively separable preferences if he/she/it assigns a nonnegative number a_{fw} to each $w \in W$ and assigns the value $v(A) =$

$\sum_{w \in A} a_{fw}$ to each allowable set A of partners for f .

Agent f compares two allowable sets by comparing the values of these sets. This concept is similarly defined for $w \in W$. If the agents have additively separable preferences, we can think that if a partnership $(f, w) \in F \times W$ is formed, then the partners participate in some joint activity that generates a payoff a_{fw} for player f and b_{fw} for player w . These numbers are fixed, i.e., they are not negotiable. If the preferences of the agents are additively separable, then they are responsive. The converse is not true (see Kraft et al. 1959).

Allowable set of partners for $f \in F$ with quota $r(f)$ is a family of elements of $F \cup W$ with k distinct W -agents, $0 \leq k \leq r(f)$, and $r(f) - k$ repetitions of f .

Continuous two-sided matching model In this model, the structure of preferences is given by utility functions which are continuous in some money variable which varies continuously in the set of real numbers. A particular case is obtained when agents place a monetary value on each possible partner or on each possible set of partners.

Discrete two-sided matching model In the discrete two-sided matching models, agents have preferences over allowable sets of partners. The allowable sets of partners for f of the type $\{w, f, \dots, f\}$ are identified with the individual agent $w \in W$ and the allowable set of partners $\{f, \dots, f\}$ is identified with f . Under this identification, agent w is *acceptable* to agent f if and only if f likes w as well as himself/herself/itself. Similar definitions and identifications apply to an agent $w \in W$. These preferences are transitive and complete, so they can be represented by ordered lists of preferences. The model can then be described by (F, W, P, r, s) , where P is the profile of preferences and r and s are the arrays of quotas for the F -agents and W -agents, respectively.

Feasible assignment for a two-sided matching model with sides F and W is an $m \times n$ matrix $x = (x_{fw})$ whose entries are zeros or ones such that $\sum_f x_{fw} \leq s(w)$ for all $w \in W$ and $\sum_w x_{fw} \leq r(f)$ for all $f \in F$. We say that

$x_{fw} = 1$ if f and w form a partnership and $x_{fw} = 0$ otherwise. A feasible assignment x corresponds to a matching μ which matches f to w if and only if $x_{fw} = 1$. Thus, if $\sum_f x_{fw} = 0$, then w is unassigned at x or, equivalently, unmatched at μ , and if $\sum_w x_{fw} = 0$, then f is likewise unassigned at x or, equivalently, unmatched at μ .

F-optimal stable matching (respectively, payoff) for a discrete (respectively, continuous) two-sided matching model is the stable matching (respectively, payoff) which is weakly preferred by every agent in F . Similarly we define the W -optimal stable matching (respectively, payoff).

Hybrid two-sided matching model is a unification of the discrete and the continuous models. It is obtained by allowing the agents of both markets to trade with each other in the same market.

Lattice property A set L endowed with a partial order relation $\geq$ has the lattice property if $\sup\{x, y\} \equiv x \vee y$ and $\inf\{x, y\} \equiv x \wedge y$ are in L , for all $x, y \in L$. The lattice is complete if all its subsets have a supremum and an infimum (see Birkhoff 1973).

Manipulable mechanism A mechanism h is *manipulable* or it is *not strategy proof* if in some revelation game induced by h , stating the true preferences is not a dominant strategy for at least one player. A mechanism h is *collectively manipulable* if in some revelation game induced by h , there is a coalition whose members can be better off by misrepresenting its preferences.

Matching mechanism For the discrete two-sided matching models, a matching mechanism is a function h whose range is the set of all possible inputs $X = (F, W, P, r, s)$ and whose output $h(X)$ is a matching for X .

Matching μ in a two-sided matching model with sides F and W is a function that maps every agent into an allowable set of partners for him/her/it, such that f is in $\mu(w)$ if and only if w is in $\mu(f)$, for every $(f, w) \in F \times W$. If we relax this condition, the function is called a *pre-matching*. A matching describes the set of partnerships of the type (f, w) , (f, f) , or (w, w) , with $f \in F$ and $w \in W$, formed by the agents. We

say that a player that does not enter any partnership is unmatched. Agents compare two matchings by comparing the two allowable sets of partners they obtain.

Maximin preferences in a discrete two-sided matching model with sides F and W . Agent $f \in F$ with a quota of $r(f)$ has a maximin preference relation over allowable sets of partners if whenever two allowable sets C and C' contained in W , such that f prefers C' to C and no w in C is unacceptable to f , then a) all of C' are acceptable to f and b) if $|C| = r(f)$, then the least preferred worker in $C' - C$ is preferred by f to the least preferred worker in $C - C'$. Similarly we define maximin preference for $w \in W$.

Choice set of $f \in F$ from $A \subseteq W(Ch_f(A))$ in a discrete two-sided matching model with sides F and W . Let $B = \{A' \mid A' \text{ is an allowable set of partners for } f \text{ and } A' \cap W \text{ is contained in } A\}$. Then, $A' \in Ch_f(A)$ if and only if $A' \in B$ and f likes A' at least as well as A'' , for all $A'' \in B$. Similarly we define $Ch_w(A)$ for $w \in W$ and $A \subseteq F$.

Outcome For the discrete two-sided matching models, the outcome is a matching or at least corresponds to a matching; for the continuous two-sided matching models, the outcome specifies a payoff for each agent and a matching.

Pareto-optimal matching A feasible matching μ is Pareto optimal if there is no feasible matching which is weakly preferred to μ by all players and it is strictly preferred by at least one of them.

Quota of an agent Quota of an agent in a two-sided matching model is the maximum number of partnerships an agent is allowed to form. When every participant can form one partnership at most, the matching model is called one to one. If only the players of one of the sides can form more than one partnership, the matching model is said to be many to one. Otherwise the matching model is many to many.

$r(f)$ -separable preference in a discrete two-sided matching model with sides F and W . Agent $f \in F$ with a quota of $r(f)$ has an $r(f)$ -separable preference relation over allowable sets of partners if whenever $A = B \cup \{w\} \setminus \{f\}$ with $w \notin B$ and $f \in B$, then f prefers A to B if and only if f prefers w to f . Similarly we define $s(w)$ -separable preference for $w \in W$.

Rematching proof equilibrium is a Nash equilibrium profile from which no pair of players $(f, w) \in F \times W$ can profitably deviate given that the other players do not change their strategies.

Responsive preference in a discrete two-sided matching model with sides F and W . Agent $f \in F$ has a responsive preference relation over allowable sets of partners if whenever (i) A and B are two allowable sets of partners for player f , (ii) j and k are two elements of $W \cup \{f\}$, and (iii) $A = B \cup \{w\} \setminus \{w'\}$ with $w \notin B$ and $w' \in B$, then f prefers A to B if and only if f prefers w to w' . Similarly we define the responsive preference for $w \in W$.

Revelation game It is the strategic game induced by a revelation mechanism for the discrete two-sided matching (F, W, P, r, s) : the set of players is given by the union of F and W ; a strategy of player j is any possible list of preferences $Q(j)$ that player j can state; the outcome function is given by the mechanism h and the preferences of the players over the set of outcomes are determined by P .

Revelation mechanism Given the discrete two-sided matching model (F, W, P, r, s) , a revelation mechanism is the restriction of a matching mechanism h to the set of discrete two-sided matching markets (F, W, Q, r, s) where the sets of agents and quotas are fixed.

Sincere strategy for a player j in a revelation game is the true list of preferences $P(j)$.

Stable matching mechanism It is a matching mechanism h such that $h(X)$ is always stable for the market X . If $h(X)$ always produces the F -optimal stable matching for X , then it is called an *F-optimal stable matching mechanism* and so on.

Stable outcome It is the natural solution concept for a two-sided matching model. It is also referred as setwise-stable outcome. See the definition below.

Strongly substitutable preferences in a discrete two-sided matching model with sides F and W . Agent $f \in F$ has a strongly substitutable preference relation over allowable sets of partners if for every pair of allowable sets of partners A and B such that $A \succ_f B$, if $w \in Ch_f(W \cap A \cup \{w\})$, then $w \in Ch_f(W \cap B \cup \{w\})$. This is a stronger condition than substitutability and responsiveness.

Substitutable preferences in a discrete two-sided matching model with sides F and W . Agent $f \in F$ has a substitutable preference relation over allowable sets of partners if whenever $A \subseteq W$ and $B \subseteq W$ are such that $A \cap B = \emptyset$, then (i) for all $S' \in Ch_f(A \cup B)$, there is some $S \in Ch_f(A)$ such that $S' \cap A \subseteq S$ and (ii) for all $S \in Ch_f(A)$, there is some $S' \in Ch_f(A \cup B)$ such that $S' \cap A \subseteq S$. If an agent's preference is responsive, then it is substitutable. When preferences are strict, conditions (i) and (ii) are equivalent to requiring that if $Ch_f(A \cup B) = S'$, then $S' \cap A \subseteq Ch_f(A)$. This concept is similarly defined for $w \in W$.

T-map in a discrete two-sided matching model with sides F and W is defined as follows. For every pre-matching μ , let $T(\mu(f)) = Ch_f(U(f, \mu))$ for all $f \in F$, where $U(f, \mu) = \{w \in W \mid f \in Ch_w(\mu(w) \cup \{f\})\}$. Similarly, $T(\mu(w)) = Ch_w(U(w, \mu))$ for all $w \in W$, where $U(w, \mu) = \{f \in F \mid w \in Ch_f(\mu(f) \cup \{w\})\}$.

Truncation Let $P(a)$ be the a 's preference list over individuals for a discrete two-sided matching model. For $a \in F \cup W$, the list of preferences $Q(a)$ over individuals is a truncation of $P(a)$ at some agent b if $Q(a)$ keeps the same ordering as $P(a)$ but ranks as unacceptable all agents which are ranked below b .

Two-sided matching model Two-sided matching model is a game theoretical model whose elements are (i) two disjoint and finite sets of agents, F with m elements and W with n elements, referred to as the sides of the matching model, (ii) the structure of agents' preferences, and (iii) the agents' quotas. The rules of the game determine the feasible outcomes. The main activity of the agents from one set is to form partnerships with the agents on the other set. Players derive their payoffs from the set of partnerships they form. The agents belonging to F and W are called F -agents and W -agents, respectively.

Definition of the Subject

This entry describes the basic elements of the cooperative and noncooperative approaches for two-sided matching models and analyzes the

fundamental differences and similarities between some of these models.

Basic Definitions

Feasible outcome	is an outcome that is specified by the rules of the game. In the discrete case, a feasible outcome is a feasible matching μ or at least corresponds to a feasible matching. The usual definition is the following. <i>The matching μ is feasible if it matches every agent to an allowable set of partners and $\mu(f) \in Ch_f(\mu(f) \cap W)$ and $\mu(w) \in Ch_w(\mu(w) \cap F)$ for every $(f, w) \in F \times W$. Then, if preferences are responsive, every matched pair is mutually acceptable. An implication of this definition is that a feasible outcome is always individually rational.</i> In the continuous case, the rules of the game may specify, for example, whether the agents negotiate their payoffs individually within each partnership or if they negotiate in blocks. In the former case, a feasible outcome specifies an array of individual payoffs for each player, indexed according to the partnerships formed under the corresponding matching. In the latter case, the feasible outcome only specifies a total payoff for each agent.
Corewise stability	is a solution concept that assigns to each two-sided matching model the set of feasible outcomes which are not <i>dominated</i> by any feasible outcome via a coalition. The feasible outcome x <i>dominates</i> the feasible outcome y via the coalition $S \neq \emptyset$, if (i) every player in S prefers x to y and (ii) if $j \in S$, then all of j's partners under x belong to S. Coalition S is said to <i>block</i> the outcome y. It is a natural solution concept for the one-to-one matching models and for the continuous many-to-one matching models.
Strong corewise stability	is a solution concept that assigns to each two-sided matching model the set of feasible outcomes which are not <i>weakly dominated</i> by any feasible outcome via a coalition. The feasible outcome x <i>weakly dominates</i> the feasible outcome y via the coalition $S \neq \emptyset$, if (i) all players in S weakly prefer x to y and at least one

(continued)

	of them strictly prefers x to y and (ii) if $j \in S$, then all of j's partners under x belong to S. Coalition S is said to <i>weakly block</i> the outcome y. Strong corewise stability is a natural solution concept for the discrete many-to-one matching models.
Pairwise stability	is a solution concept that assigns to each two-sided matching model the set of feasible outcomes which are not <i>quasi-dominated</i> by any feasible outcome via a pair of agents $(f, w) \in F \times W$. For the discrete case, an outcome can be identified with a matching. We say that the feasible matching x <i>quasi-dominates</i> the feasible matching y via the coalition $S \neq \emptyset$, if (i) every player in S prefers the matching x to the matching y and (ii) if $j \in S$ and $k \in x(j)$, then $k \in S \cup y(j)$. For the continuous case in which agents negotiate their payoffs individually with each partner, the feasible outcome x <i>quasi-dominates</i> the feasible outcome y via the coalition $S \neq \emptyset$, if (i) every player in S gets a higher total payoff under x than under y, and (ii) if $j \in S$ and $k \in x(j)$, then $k \in S$ or $k \in y(j)$, in which case, k keeps the same individual payment obtained with j under y. For the continuous case in which agents cannot negotiate their individual payments, the definition of <i>quasi-dominance</i> is equivalent to that of dominance. In any case, coalition S is said to <i>destabilize</i> the outcome y. Thus, for the discrete models, for example, a matching μ is pairwise stable if it is feasible and it is not destabilized by any pair $(f, w) \in F \times W$. The pair of agents (f, w) destabilizes the matching μ if these agents prefer each other to some of their current mates. Pairwise stability is a natural solution concept for the marriage and the college admission models as well as for continuous matching models in which the agents negotiate their individual payoffs.
Setwise stability	is the solution concept that assigns to each two-sided matching model the set of feasible outcomes which are not <i>quasi-dominated</i> by any feasible outcome via a coalition. It is a generalization of the group stability concept defined by Roth (1985b) for the college admission

(continued)

	model. An attempt to extend the concept of group stability to a discrete many-to-many matching market with substitutable and strict preferences was not successful in Roth (1984b). The concept introduced in that paper is equivalent to pairwise stability. It turns out that pairwise-stable matchings may be blocked by coalitions of size greater than two in this model. In fact, an example of a pairwise-stable matching that is corewise unstable is presented in Blair (1988) and another one in Sotomayor (1999b). Thus, the pairwise stability concept cannot be regarded as the natural cooperative solution concept for this model.
	Setwise stability regarded as a new cooperative equilibrium concept, different from the core concept, was obtained for the first time in Sotomayor (1992) in a many-to-many matching model with additively separable utilities. In this model, the set of setwise-stable outcomes equals the set of pairwise-stable outcomes and may be smaller than the core. The setwise stability concept was introduced in Sotomayor (1999b) as a refinement of the core concept and as the natural cooperative equilibrium concept for a two-sided matching model. In the discrete many-to-many matching models, it is stronger than pairwise stability plus corewise stability (see Example 3 below). Since the essential coalitions in the marriage and in the college admission models are pairs of players made up of one agent of each side of the market, a setwise-stable matching is a pairwise-stable matching in these models. If the preferences are strict, the pairwise-stable matchings are the corewise-stable matchings in the marriage model and are the strong corewise-stable matchings in the college admission model.

A Brief Historical Account

The theory of the two-sided matching markets was born in 1962, year of the publication of the seminal paper by Gale and Shapley, who

formulated the stable matching problem for the marriage and the college admission markets. The problem of “college admissions” as described in that paper involves a set of colleges and a set of students. Each student lists in order of strict preference those colleges he/she wishes to attend and each college lists in order of strict preference those students it is willing to admit. Furthermore, each college has a quota, representing the maximum number of students it can admit. The problem is then to devise some allocation procedure of students to colleges in a way that takes account of their respective preferences. More specifically, given any set of colleges and students, together with their preferences and quotas, can one find a stable matching?

The answer to this question is affirmative.

For the existence proof, Gale and Shapley constructed a simple deferred-acceptance algorithm, which, starting from the given data, leads to a stable matching in a finite number of steps. The matching obtained in this way is the unique stable matching (there may be many) which is preferred by all students to any other such matching. For this reason, it is called optimal stable matching for the students.

In the 1962 paper, they remark that: “...even though we no longer have the symmetry of the marriage problem, we can still invert our admissions procedure to obtain the unique “college optimal” assignment. The inverted method bears some resemblance to a fraternity “crush week” it starts with each college making bids to those applicants it considers most desirable, up to its quota limit, and then the bid-for students reject all but the most attractive offer, and so on.”

In the paper mentioned above, the authors express some reservation on the possibilities of application of their algorithm. It turns out that since 1951, a mathematically equivalent algorithm was being used by the National Resident Matching Program (NRMP), located in Evanston, Illinois. The NRMP has the task each year of assigning graduates of all medical schools in the United States to hospitals where they are required to serve a year’s residency. The algorithm used by the NRMP was mathematically the equivalent to the one described in Gale and Shapley (1962) to produce the optimal stable matching for the

colleges. Thus, in this algorithm, each hospital applies to its quota of students. The confirmation of this fact was obtained by David Gale in 1975. In a letter from December 8, 1975, in response to a letter from Gale to the NRMP, Elliott Peranson, consultant to NRMP, responsible for the technical operation of the matching program, says the following:

However, I might point out that the NRMP algorithm in fact uses the inverse procedure and produces the unique “college optimal” assignment rather than this “student optimal” assignment. This procedure more closely parallels the actual admissions process where a matching algorithm is not used. In this case students apply to all hospitals they would consider (not just their first choice), each hospital then selects the most desirable students, up to its quota limit, from amongst all applicants, then the “bid-for” students reject all but the most desirable offer, and so on.

Hence, the proof that the NRMP was yielding a stable matching, which was the optimal stable matching for the colleges, is that such an outcome is always obtained by the Gale and Shapley's algorithm with the colleges proposing.

The discovery that the two algorithms were mathematically equivalent was first spread orally and later reported in Gale and Sotomayor (1983) with an equivalent description of the NRMP algorithm (Roth (1984c) also presents the NRMP algorithm). This was the first application of the matching theory of which we have knowledge. The algorithms proposed by Gale and Shapley are described below. Their description is quoted from Gale and Sotomayor (1983).

Gale-Shapley Algorithm with the Colleges Proposing to the Applicants

Each hospital H tentatively admits its quota q_H consisting of the top q_H applicants on its list. Applicants who are tentatively admitted to more than one hospital *tentatively accept* the one they prefer. Their names are then removed from the lists of all other hospitals which have tentatively admitted them. This gives the *first tentative matching*. Hospitals which now fall short of their quota again admit tentatively until either their

quota is again filled or they have exhausted their list. Admitted applicants again reject all but their favorite hospital, giving the *second tentative matching*, etc. The algorithm terminates when, after some tentative matching, no hospitals can admit any more applicants either because their quota is full or they have exhausted their list. The tentative matching then becomes permanent.

Gale-Shapley Algorithm with the Applicants Proposing to the Colleges

Each applicant *petitions* for admission to his/her favorite hospital. In general, some hospitals will have more petitioners than allowed by their quota. Such oversubscribed hospitals now reject the lowest petitioners on their preference list so as to come within their quota. This is the *first tentative matching*. Next, rejected applicants petition for admission to their second favorite hospital and again oversubscribed hospitals reject the overflow, etc. The algorithm terminates when every applicant is tentatively admitted or has been rejected by every hospital on his/her list.

The fact that the matching produced by the NRMP algorithm is stable stands for one of the most important applications of game theory to economics. During about fifty years, the allocation procedures used to assign interns to hospitals in the United States produced unstable matchings. Unsuccessful procedures were often proposed by the Association of American Medical Colleges. This sort of events culminated with a centralized mechanism that employed the NRMP algorithm. Such a centralized mechanism lasted for almost fifty years, suggesting that interns and hospitals had reached an equilibrium. And the paper written by Gale and Shapley corroborated that the game theoretical predictions were, once more, correct.

Before the publication of Gale and Sotomayor (1983), a few, but important, contributions were made to the theory of two-sided matching markets. The famous paper of 1972 by Shapley and Shubik establishes the assignment game via the introduction of money, as a continuous variable, into the marriage model. The book *Marriage Stables* by Knuth was published in 1976. In this volume, the proof,

attributed to Conway, that the set of stable matchings for the marriage model is a lattice is presented. The assignment game was generalized by Kelso and Crawford (1982) to a model where the utilities satisfy some gross substitute condition. Another generalization, which considers continuous utility functions, non-necessarily linear, was presented in Demange and Gale (1985).

Nevertheless, among the contributions of this period, one of them caused considerable impact. This was the non-manipulability theorem by Dubins and Freedman (1981). In a stable revelation mechanism, for every profile of preferences that can be selected by the agents, some algorithm that yields a stable matching is used. These authors prove that the revelation mechanism that produces the optimal stable matching for a given side of the marriage market is not collectively manipulable by the agents of that side. Also, this non-manipulability result holds for the college admission market when the mechanism yields the student-optimal stable matching. An analog to this result was proved in Demange and Gale (1985) for the continuous model through a key lemma that became known in the literature as the blocking lemma. The main challenge that motivated Gale and Sotomayor (1983) was to prove the discrete version of the blocking lemma, which allowed to prove the non-manipulability theorem in just three lines. Two simple and short proofs (one with the use of the algorithm and the other one without the use of the algorithm) of Dubins and Freedman's theorem were presented as an alternative to the original proof by the authors which was about twenty pages long. An example in Dubins and Freedman (1981), where some woman can be better off by falsifying her preference list when the man-optimal stable matching is to be employed, motivated Gale and Sotomayor (1985). This paper proves that such a mechanism can almost always be manipulated by the women and then treats the strategic possibilities for these agents in the corresponding strategic game. Another paper of this period was Roth (1982), which proves, via an example, that any rule for selecting a stable matching is manipulable (either by some man or some woman).

The existence theorem of manipulability by the women, the impossibility theorem, and the non-

manipulability theorem attracted the authors' interest toward a fruitful line of investigation concerning the incentives facing the agents when an allocation mechanism is employed. Algorithms have been developed for this purpose for several matching models. The games induced by such mechanisms are played noncooperatively by the agents, and in general, their self-enforcing agreements lead to a stable outcome. In these cases, a noncooperative implementation of the set of stable outcomes is provided. Analyzing the strategic behavior of the agents in such games has been an important subject of research of several authors in an attempt to get precise answers to the strategic questions raised. In this direction, we can cite Roth (1982; Roth 1984a), Gale and Sotomayor (1985), Perez-Castrillo and Sotomayor (2002), Sotomayor (2004a, b, c), Kamecke (1989), Kara and Sönmez (1996, 1997), Alcalde (1996), and Alcalde et al. (1998), among others.

Over all these years, the popularity of the matching theory has spread among mathematicians and economists, mainly due to the publication in 1990 of the first edition of the book *Two-Sided Matchings: A Study in Game Theoretic Modeling and Analysis*, by Roth and Sotomayor, which attempts a comprehensive survey of the main results on the two-sided matching theory until that date (an extensive bibliography can also be found in <http://kuznets.fas.harvard.edu/~aroth/bib.html#matchbib>).

The stable matching problem has been generalized to several two-sided matching models, which have been widely modeled and analyzed under both cooperative and noncooperative game-theoretic approaches. Through these models, a variety of markets has become better understood, which has considerably contributed to their organization. The deferred acceptance algorithm of Gale and Shapley and adaptations of it have been applied in the reorganization of admissions processes of many two-sided matching markets. And the design of these mechanisms has also raised new theoretical questions (in this connection, see, e.g., Balinski and Sönmez (1999), Ergin and Sönmez (2006), and Pathak and Sönmez (2006), Abdulkadiroglu and Sönmez (2003), and Bardella and Sotomayor (2006)).

Introduction

The two-sided matching theory is the study of game theoretical models in which the set of players is partitioned into the disjoint union of two finite sets and the main activity of the agents from one set is to form partnerships with the agents on the other set. In addition, a structure of preferences is available for the players, as well as an array of quotas, one quota for each participant, representing the maximum number of partnerships that he/she/it is allowed to form.

The rules of the game determine the feasible outcomes. These models are called two-sided matching models. They are suitable for modeling a great variety of labor markets of firms and workers, markets of buyers and sellers, markets of students and schools, etc. They are grouped into three categories: the discrete, the continuous, and the hybrid matching models. The distinctions between them are based on the kind of preference structure they are endowed with. Within each category, the models vary as to the possibilities of the agents to form multiple partnerships, the rules of the game, and some variation of the structure of the preferences.

Given the sets of agents, the structure of preferences and quotas, and the rules of the game, the question that emerges is:

1. Which partnerships should be formed?
2. If a partnership is formed, what payoff should be awarded to each agent?

The prediction should be outcomes that cannot be upset by any coalition. The idea is that, in games where players form partnerships, it should be allowed for coalitions to be formed in which its members keep some of their current partners if they wish and replace some others with new partners belonging to the coalition. Furthermore, by doing this, they get a preferable outcome. This intuitive idea is captured by a refinement of the core concept introduced in Sotomayor (1999b), called setwise stability (stability, for short). For the marriage model and the college admission model with responsive preferences, this concept coincides with the one introduced by Gale and Shapley.

This review is an attempt to understand some of the differences and similarities between some matching models. We do that by analyzing both the cooperative and the noncooperative structure of these models. Some future directions will then be presented.

The rest of the article is organized as follows. Section “[Discrete Two-Sided Matching Models](#)” is devoted to the cooperative approach of the discrete matching models with responsive preferences. It introduces the basic cooperative one-to-one, many-to-one, and many-to-many matching models and discusses the main properties that characterize the set of stable outcomes for these models. This kind of analysis is also provided in section “[Continuous Two-Sided Matching Model with Additively Separable Utility Functions](#)” for the continuous matching models where the utilities are additively separable. Section “[Hybrid One-to-One Matching Model](#)” discusses some fundamental similarities and differences between a one-to-one hybrid model and its corresponding non-hybrid matching models. Strategic questions are treated in section “[Incentives](#).” Section “[Future Directions](#)” presents some future directions and open problems.

Discrete Two-Sided Matching Models

There are two finite and disjoint sets of agents, F and W , which we may think of as being sets of men and women, colleges and students, firms and workers, etc. To fix ideas, let us describe the model in terms of firms and workers. Then, F is a set of m firms and W is a set of n workers. Salaries cannot be negotiated and are part of the job description. Each worker w has a quota $s(w)$ representing the maximum number of jobs in different firms he/she can take. Each firm f has a quota $r(f)$ representing the maximum number of workers it can hire. Set r and s , the array of quotas of the firms and the workers, respectively.

Since the agents form partnerships, they always have preferences over potential individual partners, that is, over allowable sets of partners with only one agent belonging to the opposite side of the market. These preferences are assumed to be strict. They are transitive and complete and so,

they can be represented by ordered lists of preferences. Thus, the individual preference relation of firm f can be represented by an ordered list of preferences $P(f)$ on the set $W \cup \{f\}$; the individual preference relation of worker w can be represented by an ordered list of preferences $P(w)$, on the set $F \cup \{w\}$. The array of these preferences will be denoted by P . Then, an agent w is acceptable to an agent f if $w \geq_f f$. Similarly, an agent f is acceptable to an agent w if $f \geq_w w$.

If an agent may form more than one partnership, then his/her/its preferences are not restricted to the individual potential partners. That is, agents have preferences over any allowable set of partners. Given two allowable sets of partners, A and B , for agent $y \in F \cup W$, we write $A >_y B$ to mean y prefers A to B and $A \geq_y B$ to mean y likes A at least as well as B . In order to fix ideas, we will assume that these preferences are *responsive* to the agents' individual preferences and are not necessarily strict.

The rules of the market are that any firm and worker pair may sign an employment contract with each other if they both agree; any firm may choose to keep one or more of its positions unfilled, and any worker may choose not to fill his or her quota of jobs if he/she wishes to do so.

When the quota of every agent is one, we have the *marriage model*. In this case, an allowable set of partners for any agent is a singleton. Therefore, every agent only has preferences over individuals.

If only the agents of one of the sides are allowed to have a quota greater than one, then we have the so-called *college admission model* with responsive preferences.

In both models, it is a matter of verification that the sets of setwise-stable matchings, pairwise-stable matchings, and strong corewise-stable matchings coincide. For the many-to-many case, this is not always true. The strong corewise stability concept is not a natural solution concept for this model as Example 1 shows.

Example 1 (Sotomayor 1999b) (The corewise stability concept is not a natural solution concept for the many-to-many case) Consider two firms f_1 and f_2 and two workers w_1 and w_2 . Each firm may employ and wants to employ both

workers; worker w_1 may take, at most, one job and prefers f_1 to f_2 ; worker w_2 may work and wants to work for both firms. If the agents can communicate with each other, the outcome that we expect to observe in this market is obvious: f_1 hires both workers and f_2 hires only worker w_2 . Of course, this outcome is in the strong core. Since f_1 has a quota of two and w_1 prefers f_1 to f_2 , we cannot expect to observe the strong corewise-stable outcome where f_1 hires only w_2 and f_2 hires both workers. That is, both outcomes are in the strong core, but only the first one is expected to occur. Our explanation for this is that only the first outcome is setwise stable.

On the other hand, the pairwise stability concept is not a refinement of the core for the discrete many-to-many matching models. See Example 2.

Example 2 (Sotomayor 1999b) (A pairwise-stable matching which is not in the core) Consider the following labor market with a set of firms $F = \{f_1, f_2, f_3, f_4\}$ and a set of workers $W = \{w_1, w_2, w_3, w_4\}$, where each firm can hire two workers and each worker can work for two firms. If firm f_i hires worker w_j , then f_i gets the profit a_{ij} and w_j gets the salary b_{ij} . The pairs of numbers (a_{ij}, b_{ij}) are given in Table 1.

Consider the matching μ where f_1 and f_2 are matched to $\{w_3, w_4\}$ and f_3 and f_4 are matched to $\{w_1, w_2\}$ (the payoffs of each matched pair are presented in boldface). This matching is pairwise stable. In fact, f_3 and f_4 do not belong to any pair which causes instability, because they are matched to their two best choices: w_1 and w_2 ; (f_1, w_1) and (f_1, w_2) do not cause instabilities since f_1 is the worst choice for w_1 and w_2 is the worst choice for f_1 ; (f_2, w_1) and (f_2, w_2) do not cause instabilities since w_1 is the worst choice for f_2 and f_2 is the worst choice for w_2 .

Two-Sided Matching Models, Table 1 Payoff matrix of Example 2. Each row represents a firm and each column a worker. The values in the cell represent the payoffs to the corresponding firm (*first value*) and worker (*second value*)

10,1	1,10	4,10	2,10
1,10	10,1	4,4	2,4
10,4	4,4	2,2	1,2
10,2	4,2	2,1	1,1

Nevertheless, f_1 and f_2 prefer $\{w_1, w_2\}$ to $\{w_3, w_4\}$ and w_1 and w_2 prefer $\{f_1, f_2\}$ to $\{f_3, f_4\}$. Hence this matching is not in the core, since it is blocked by the coalition $\{f_1, f_2, w_1, w_2\}$.

Example 3 shows that setwise stability is a strictly stronger requirement than pairwise stability plus strong corewise stability. It presents a situation in which the set of setwise-stable matchings is a proper subset of the intersection of the strong core with the set of pairwise-stable matchings.

Example 3 (Sotomayor 1999b) (A strong corewise-stable matching which is pairwise stable and is not setwise stable) Consider the following labor market with a set of firms $F = \{f_1, f_2, f_3, f_4, f_5, f_6\}$ and a set of workers $W = \{w_1, w_2, w_3, w_4, w_5, w_6, w_7\}$, where $r_1 = 3, r_2 = r_5 = 2, r_3 = r_4 = r_6 = 1, s_1 = s_2 = s_4 = 2$, and $s_3 = s_5 = s_6 = s_7 = 1$. If firm f_i hires worker w_j , then f_i gets profit a_{ij} and the worker gets a salary b_{ij} . The pairs of numbers (a_{ij}, b_{ij}) are given in Table 2.

Let μ be the matching given by

$$\begin{aligned} \mu(f_1) &= \{w_2, w_3, w_7\}, \quad \mu(f_2) = \{w_5, w_6\}, \\ \mu(f_3) &= \mu(f_4) = \{w_1\}, \quad \mu(f_5) = \{w_2, w_4\}, \quad \text{and} \\ \mu(f_6) &= \{w_4\}. \end{aligned}$$

The associated payoffs are shown in bold in Table 2. This matching is in the strong core. In fact, if there is a matching μ' which weakly dominates μ via some coalition A , then, under μ' , no player in A is worse off and at least one player in A is better off. Furthermore, matching μ' must match all members of A among themselves. By inspection, we can see that the only players that can be better off are f_1, f_2, w_1 , and w_4 , for all

remaining players are matched to their best choices. However, if A contains one player of the set $\{f_1, f_2, w_1, w_4\}$, then A must contain all four players. In fact, if $f_1 \in A$, then f_1 must form a new partnership with w_1 , so $w_1 \in A$. If $w_1 \in A$, then w_1 must form a new partnership with f_2 , so $f_2 \in A$. If $f_2 \in A$, then f_2 must form a new partnership with w_4 , so $w_4 \in A$. Finally, if $w_4 \in A$, then w_4 must form a new partnership with f_1 , so $f_1 \in A$. Thus, if μ' weakly dominates μ via A , then f_1, f_2, w_1 , and w_4 are in A and f_1 and f_2 form new partnerships with w_1 and w_4 . Nevertheless, f_1 must keep his partnership with w_2 , his best choice. Then w_2 must be in A , so she cannot be worse off and so f_5 must also be in A . But f_5 requires the partnership with w_4 , who has quota of 2 and has already filled her quota with f_1 and f_2 . Hence f_5 is worse off at μ' than at μ and then μ' cannot weakly dominate μ via A .

The matching μ is clearly pairwise stable. Nevertheless, the coalition $\{f_1, f_2, w_1, w_4\}$ causes an instability in μ . In fact, if f_1 is matched to $\{w_1, w_2, w_4\}$ and f_2 is matched to $\{w_1, w_4\}$, then f_1 gets 28 and the rest of the players in the coalition get 11 instead of 6. Hence, the matching μ is not setwise stable.

The question is then to know if, given any two sets of agents with their respective preferences and quotas, one can always find a setwise-stable matching. The answer is affirmative for the one-to-one matching model and for the many-to-one matching models with substitutable preferences. The existence of setwise-stable matchings for the marriage model was first proved by Gale and Shapley (1962). Sotomayor (1996a) also provides a simple proof of the existence of stable matchings for the marriage model that connects stability with a broader notion of stability with respect to unmatched agents. Gale and Shapley construct a deferred acceptance algorithm described below and prove that it yields a stable matching in a finite number of steps.

The deferred acceptance algorithm with the men making the proposals. Each man begins by proposing to his favorite woman (the first woman on his preference list). Each woman rejects the proposal of any man unacceptable to her, and in case she gets several proposals, she keeps only her most preferred one. If a man is not rejected at this step, he is kept engaged. At any step, any man

Two-Sided Matching Models, Table 2 Payoff matrix of Examples 3 and 4. Each row represents a firm and each column a worker. The values in the cell represent the payoffs to the corresponding firm (first value) and worker (second value)

13,1	14,10	4,10	1,10	0,0	0,0	3,10
1,10	0,0	0,0	10,1	4,10	2,10	0,0
10,4	0,0	0,0	0,0	0,0	0,0	0,0
10,2	0,0	0,0	0,0	0,0	0,0	0,0
0,0	9,9	0,0	10,4	0,0	0,0	0,0
0,0	0,0	0,0	10,2	0,0	0,0	0,0

who was rejected at the previous step proposes to his next choice (his most preferred woman among those who have not rejected him), as long as there remains an acceptable woman to whom he has not yet proposed (if at one step a man has been rejected by all of his acceptable women, he issues no further proposals). The algorithm terminates after any step in which no man is rejected, because then every man is either engaged to some woman or has been rejected by every woman on his list of acceptable women. Women who did not receive any acceptable proposals, and men rejected by all women acceptable to them remain single.

To see that the matching yielded by this algorithm is stable, first observe that no agent has an unacceptable partner. In addition, if there is some man f and woman w not matched to each other and such that f prefers w to his current partner, then woman w is acceptable to man f and so he must have proposed to her at some step of the algorithm. But then he must have been rejected by w in favor of someone she prefers to f . Therefore, w is matched to a man whom she prefers to f by the transitivity of the preferences and so f and w do not destabilize the matching.

For the college admission model with responsive preferences, Gale and Shapley defined a related marriage model in which each college is replicated a number of times equal to its quota, so that in the related model, every agent has a quota of one. If $f_1, \dots, f_{r(f)}$ are the $r(f)$ copies of college f , then each of these f_i 's has preferences over individuals that are identical with those of f . Each student's preference list is changed by replacing f , wherever it appears on his/her list, by the string $f_1, \dots, f_{r(f)}$ in that order of preference. Therefore, the stable matchings of the related marriage market are in natural one-to-one correspondence with the stable matchings of the college admission market. By using the existence theorem for the marriage model, we obtain the corresponding result for the college admission model.

The existence proof for the many-to-one case with strict and substitutable preferences was first given by Kelso and Crawford (1982) through a variant of the deferred-acceptance algorithm.

If the market does not have two sides or the many-to-one matching model does not have substitutable preferences, then the set of setwise-stable

matchings may be empty. Gale and Shapley (1962) present an example of a one-sided matching model that does not have any stable matchings. This model was called by these authors the "roommate problem." In this example, there are four agents $\{a, b, c, d\}$, such that a 's first choice is b , b 's first choice is c , c 's first choice is a , and d is the last choice of all the other agents. Of course, if some agent is unmatched, then there will be two unmatched agents and they will destabilize the matching. If every one is matched, the agent who is matched to d will form a blocking pair with the agent who lists him at the head of his list. Therefore, there is no stable matching in this example.

A small amount of literature has grown around the issues of finding conditions in which the set of stable matchings is nonempty for the roommate problem and the performance of algorithms that can produce them when they exist (see Abeledo and Isaak (1991), Irving (1985), Tan (1991), Chung (2000), and Sotomayor (2005)).

If preferences are not substitutable, Example 2.7 of Roth and Sotomayor (1990) shows that setwise-stable matchings may not exist in the many-to-one case.

Even in the simplest case of preferences representable by additively separable utility functions, setwise-stable matchings may not exist for the many-to-many case. See the example below.

Example 4 (Sotomayor 1999b) (Nonexistence of stable matchings) Consider again the matching model of Example 2. We are going to show that the set of setwise-stable matchings is empty. First, observe that f_3 prefers $\{w_1, w_2\}$ to any other set of players and f_3 is the second choice for w_1 and w_2 ; w_3 prefers $\{f_1, f_2\}$ to any other set of players and w_3 is the second choice for f_1 and f_2 . Then, in any stable matching μ , f_3 must be matched to w_1 and w_2 , while w_3 must be matched to f_1 and f_2 . Separate the cases by considering the possibilities for the second partner of w_1 , under a supposed stable matching μ :

1. (w_1 is matched to $\{f_2, f_3\}$). Then f_2 is not matched to w_4 and we have that $\{f_2, w_4\}$ causes an instability in the matching, since f_2 prefers w_4 to w_1 and f_2 is the second choice for q_4 .

2. (w_1 is matched to $\{f_3, f_4\}$). Then the following possibilities occur:
 - (a) (w_2 is matched to $\{f_3, f_4\}$.) Then $\{f_1, f_2, w_1, w_2\}$ causes an instability in the matching. This matching is pairwise stable, but it is not in the core.
 - (b) (w_2 is matched to $\{f_3, f_1\}$.) Then $\{f_1, w_4\}$ causes an instability in the matching, since f_1 is the first choice for w_4 and f_1 prefers w_4 to w_2 .
 - (c) (w_2 is matched to $\{f_3, f_2\}$ or $\{f_3\}$.) Then $\{f_4, w_2\}$ causes an instability in both cases, since w_2 is the second choice for f_4 and w_2 prefers f_4 to f_2 and prefers f_4 to have an unfilled position.
3. (w_1 is matched to $\{f_1, f_3\}$ or $\{f_3\}$.) Then $\{f_4, w_1\}$ causes an instability in both cases, since w_1 is the first choice for f_4 and w_1 prefers f_4 to f_1 and prefers f_4 to have an unfilled position.

Hence, there are no stable matchings in this example.

Pairwise-stable matchings always exist when the preferences are substitutable. When preferences are strict, Roth (1984b) presents an algorithm that finds a pairwise-stable matching for a many-to-many matching model with substitutable preferences. Sotomayor (1999b) provides a simple and non-constructive proof of the existence of pairwise-stable matchings for the general discrete many-to-many matching model with substitutable and not necessarily strict preferences. Martínez et al. (2004) construct an algorithm, which allows finding the whole set of pairwise-stable matchings, when they exist, for the many-to-one matching model.

Authors have looked for sufficient conditions on the preferences of the agents for the existence of setwise-stable matchings in the many-to-many cases. Sotomayor (2004b) proves that if the preferences of the firms satisfy the *maximin* property, then the set of pairwise-stable matchings coincides with the set of setwise-stable matchings. An example in that paper shows that the set of setwise stable matchings may be empty if this condition is not satisfied. It is assumed there that the preferences are responsive and it is conjectured that the result above extends to the case of substitutable preferences.

Echenique and Oviedo (2006) also address this problem with a different condition. They show that if agents on one side of the market have

strongly substitutable preferences, while the other side has substitutable preferences, then the set of setwise stable matchings coincides with the set of pairwise-stable matchings.

Konishi and Ünver (2006) give conditions on the preferences of the agents in a many-to-many matching market under which a pairwise-stable matching cannot be quasi-dominated by a pairwise-unstable matching via a collusion.

Eeckhout (2000), under the assumption of strict preferences and that every man (woman) is acceptable to every woman (man), presents a sufficient condition for uniqueness of the stable matchings in the marriage market. The condition on preferences is simple: for every $f_i \in F = \{f_1, f_2, \dots, f_m\}$, $w_i > f_i w_k$ for all $k > i$ and for every $w_j \in W = \{w_1, w_2, \dots, w_n\}$, $f_i > w_j f_k$ for all $k > i$.

One line of investigation that has been developed in the theory of two-sided matchings concerns the mathematical structure of the set of stable matchings, because it captures fundamental differences and similarities between the several kinds of models. For the marriage model and the college admission model with responsive preferences, assuming that the preferences over individuals are strict, the set of setwise-stable (stable for short) matchings have the following characteristic properties:

1. Let μ and μ' be stable matchings. Then $\mu \geq F\mu'$ if and only if $\mu' \geq W\mu$.

That is, there exists an opposition of interests between the two sides of the market along the whole set of stable matchings.

2. Every agent is matched to the same number of mates under every stable matching.

Consequently, if an agent is unmatched under some stable matching, then he/she/it is unmatched under any other stable matching.

When preferences are strict, there are two natural partial orders on the set of all stable matchings. The partial order $\geq F$ is defined as follows: $\mu \geq F\mu'$ if $\mu(f) \geq f\mu'(f)$ for all $f \in F$. The partial order $\geq W$ is analogously defined. The fact that these partial orders are well defined follows from A1.

3. The set of stable matchings has the algebraic structure of a complete lattice under the partial orders $\geq F$ and $\geq W$.

The lattice property means the following: If μ and μ' are two stable matchings, then some workers (respectively, firms) will get a preferable set of mates under μ than under μ' and others will be better off under μ' than under μ . The lattice property implies that there is then a stable matching which gives each agent the most preferable of the two sets of partners and also one which gives each of them the least preferred set of partners. That is, if μ and μ' are stable matchings, the lattice property implies that the functions λ , ν , η , and τ are stable matchings, where $\lambda = \mu \vee_F \mu'$ is defined by $\lambda(f) = \max\{\mu(f), \mu'(f)\}$ and $\lambda(w) = \min\{\mu(w), \mu'(w)\}$, the function $\eta = \mu \vee_W \mu'$ is analogously defined, and the function $\nu = \mu \wedge_F \mu'$ is defined by $\nu(f) = \min\{\mu(f), \mu'(f)\}$ and $\nu(w) = \max\{\mu(w), \mu'(w)\}$. Analogously we define $\tau = \mu \wedge_W \mu'$ (notice that $\mu \vee_F \mu'$ is the same as $\mu \wedge_W \mu'$ and $\mu \vee_W \mu'$ is the same as $\mu \wedge_F \mu'$).

The fact that the lattice is complete implies the existence and uniqueness of a maximal element and a minimal element in the set of stable payoffs, with respect to the partial order that is being considered. Thus, there exists one and only one stable matching μ_F and one and only one stable matching μ_W such that $\mu_F \geq_F \mu$ and $\mu_W \geq_W \mu$ for all stable matchings μ . Property A1 then implies that $\mu \geq_W \mu_F$ and $\mu \geq_F \mu_W$. That is:

1. *There is an F-optimal stable matching μ_F with the property that for any stable matching μ , $\mu_F \geq_F \mu$ and $\mu \geq_W \mu_F$; there is a W-optimal stable matching μ_W with symmetrical properties.*

Property A1 was first proved by Knuth (1976) for the marriage model. The result for the college admission model with responsive preferences was proved in Roth and Sotomayor (1990) by making use of the following proposition of Roth and Sotomayor (1989): *Suppose colleges and students have strict individual preferences, and let μ_1 and μ_2 be stable matchings for the college admission model such that $\mu_1(f) \neq \mu_2(f)$. Let μ_1^* and μ_2^* be the stable matchings corresponding to μ_1 and μ_2 in the related marriage model. If $\mu_1^*(f_i) >_f \mu_2^*(f_i)$ for some position f_i of f then $\mu_1^*(f_j) \geq \mu_2^*(f_j)$ for all positions f_j of f .*

Property A2 was proved by Gale and Sotomayor (1983) for both models. For the college admission model with responsive preferences, Roth (Roth 1986) added that *if a college does not fill its quota at some stable matching, then it has the same set of mates at every stable matching*. The restriction of property A2 to the marriage model where every pair of partners is mutually acceptable was proved by McVitie and Wilson (1970).

For the many-to-one case with substitutable preferences, Martínez et al. (2001) presents an example in which there are agents who are unmatched under some stable matching and are matched under another one. By introducing quotas in the model with substitutable preferences of Roth and Sotomayor (1990), these authors prove that if the preferences of the colleges are strict, substitutable, and $r(f)$ separable for every college f , then property A2 holds. Furthermore Roth's result mentioned above also applies.

The lattice property of the set of stable matchings for the marriage model is attributed by Knuth (1976) to Conway. The existence of the optimal stable matchings for each side of the marriage market and the college admission market with responsive preferences was first proved in Gale and Shapley (1962) by using the deferred acceptance procedure. The idea of their elegant proof is to show that a proposer is never rejected by an achievable mate, so he/she/it ends up with his/hers/its best achievable mate.

The lattice property for the college admission model with responsive preferences was obtained in Roth and Sotomayor (1990). To show that the functions λ , ν , η , and τ above are well defined, these authors used the following theorem from Roth and Sotomayor (1989): *If colleges and students have strict preferences over individuals, then colleges have strict preferences over those groups of students that they may be assigned at stable matchings. That is, if μ_1 and μ_2 are stable matchings, then a college f is indifferent between $\mu_1(f)$ and $\mu_2(f)$ only if $\mu_1(f) = \mu_2(f)$.*

This result is an immediate consequence of the proposition mentioned above, due to the responsiveness of the preferences. Therefore, if μ_1 and μ_2 are two stable matchings, then f prefers $\mu_1(f)$ to $\mu_2(f)$ if and only if the $r(f)$ most preferred students by f in the set formed by the union of $\mu_1(f)$ and $\mu_2(f)$ are those ones in $\mu_1(f)$.

For the many-to-many matching market with strict and substitutable preferences, Blair (1988) proved that the set of pairwise-stable matchings (not necessarily setwise stable) has the lattice structure under some partial order relation that is not defined by the preferences of the agents. The definition of the partial order $\geq F$ uses that if μ_1 and μ_2 are pairwise-stable matchings, then $\mu_1 \geq_F \mu_2$ if and only if $Ch(\mu_1(f) \cup \mu_2(f)) = \mu_1(f)$, for every agent $f \in F$. Similarly the partial order $\geq W$ is defined. As remarked above, the partial order defined by Blair coincides with the one defined by the preferences of the players in the college admission model with responsive preferences.

Adachi (2000) introduces a map, which is called a T -map, defined over the set of pre-matchings, in order to show that *the set of stable matchings is a nonempty lattice in the marriage market under strict preferences*. Adachi defines the T -map as follows: given a pre-matching μ , $T(\mu(f))$ is f 's most preferred worker in $\{w \in W | f \geq_w \mu(w)\} \cup (f)$ for all $f \in F$, and similarly, $T(\mu(w))$ is w 's most preferred firm in $\{f \in F | w \geq_f \mu(f)\} \cup (w)$ for all $w \in W$. Clearly, any fixed point of the T -map is a matching, and Adachi showed it has to be stable. Using the partial order $\geq F$ defined by the agents' preferences, he showed that the set of pre-matchings endowed with this partial order is a complete lattice and the T -map is an isotone function (order preserving). Thus, Tarski's fixed point theorem implies that the set of fixed points of the T -map, which is the set of stable matchings, is a nonempty complete lattice.

Theorem 1 (Tarski's Theorem (Tarski 1955))

Let E be a complete lattice with respect to some partial order $\geq$, and let f be an isotone function from E to E . Then the set of fixed points of f is nonempty and is itself a complete lattice with respect to the partial order $\geq$.

In this same vein, Echenique and Oviedo (2004, 2006) extend Adachi's (2000) approach and the T -map in order to analyze the many-to-one and many-to-many models, respectively. Again, any fixed point of the T -map is a matching. Echenique and Oviedo (2004) show that for the many-to-one model, the set of fixed points of the

T -map is equal to the set of stable matchings. By making successive iterations of the T -map, starting from some specific pre-matching, until a fixed point is reached, this map can be used to find all the stable matchings, as long as they exist. This procedure is called the T algorithm. These authors show that *as long as the strong core is nonempty, the T algorithm always converges, and if the strong core is empty, it cycles*. They present an example of a situation in which the preferences are not substitutable and the T algorithm finds strong core allocations, but the algorithm with firms proposing according to their preference lists over allowable sets of workers does not do so. Finally, they give a bound on the computational complexity of the T algorithm and show how it can be used to calculate both the supremum and infimum under Blair's partial order, which, under non-substitutable preferences, might not be easily computed. Furthermore, under substitutability, the set of pre-matchings endowed with the partial order defined by Blair is again a complete lattice and the T -map is isotone, so Tarski's theorem implies that *the strong core is a nonempty lattice under the partial order introduced in Blair (1988)*.

For the many-to-many model, the set of fixed-points of the T -map studied by Echenique and Oviedo (2006) is shown to be, under substitutability, equal to the set of pairwise-stable matchings and a superset of the set of setwise-stable matchings. Furthermore, the set of pre-matchings endowed with Blair's partial order is again a complete lattice. Then Tarski's theorem applies and the set of pairwise-stable matchings is a nonempty complete lattice. *If both sides of the market satisfy the strong substitutability property, then the set of setwise stable matchings is a complete lattice, both for Blair's partial order and for the partial order defined by the agents' preferences.*

Martínez et al. (2004) propose an algorithm which allows them to calculate the whole set of pairwise-stable matchings under substitutability.

Echenique and Yenmez (2007) study the college admission problem when students have preferences over colleagues. Using the T -map, they construct an algorithm, which finds all the core allocations, as long as the core is nonempty. In a

similar setup, Pycia (2007) finds necessary and sufficient conditions for the existence of stable matchings. Dutta and Massó (1997) studied a many-to-one version of the model of Echenique and Yenmez (2007). They showed that under certain conditions on preferences, the core is nonempty.

Hatfield and Milgrom (2005) present a general many-to-one model in which, the central concept is that of a contract, which allows a different formalization of a matching. In their model, there are a finite set of firms, a finite set of workers and a finite set of wages offered by firms. Each contract c specifies the pair (f, w) involved and the wage the worker w gets from firm f , so that the set of contracts is $C = F \times W \times WP$ (where WP is the set of wages offered by the firms). Clearly, if each firm offers a unique wage level to all workers, their model is a college admission model. Agents have preferences over the contracts in which they could be involved. In this model, a feasible allocation is a set of contracts $C' \subseteq C$ in which each worker w appears at most in one contract $c \in C'$ and each firm f appears in at most $r(f)$ contracts $c_1, \dots, c_{r(f)} \in C'$ and such that for each agent $a \in F \cup W$, we have that $Ch_a(C_a) = C_a$, where C_a is the set of contracts in C' in which agent a appears. According to Hatfield and Milgrom, a feasible allocation C' is stable if there does not exist an alternative feasible allocation which is strictly preferred by some firm f and weakly preferred by all of the workers it hires. Making use of Tarski's fixed point theorem, they prove that if preferences are strict and satisfy substitutability over the set of contracts, then the set of stable allocations is a nonempty lattice. They introduce the condition of the law of aggregate demand on preferences, which requires that for all allowable sets X, Y , if $X \subseteq Y$, then $|Ch_f(Y)| \geq |Ch_f(X)|$. By assuming that firms' preferences satisfy substitutability and the law of aggregate demand, they prove that some characteristic results on the structure of the set of stable allocations and analyze the incentives facing the agents when a mechanism which produces the W -optimal stable allocation is adopted. They show that under this mechanism, it is a dominant strategy for the workers to state their true preferences.

Ostrovsky (2008) generalizes the model presented in Hatfield and Milgrom (2005) to a K -sided discrete many-to-many matching model. A set of contracts, which allows the production and consumption of some goods, is called a network. This author considers supply networks where goods are sold and transformed through many stages, starting from the suppliers of initial inputs, going through different intermediaries until they reach the final consumer. He generalizes the concept of pairwise stability to this setting and calls it *chain stability*, which requires that there does not exist a chain of contracts such that all members of this chain are better off. A chain of contracts specifies a sequence of agents, each of whom is the seller in one contract and the buyer in the next one. Under certain conditions on the preferences of the agents, he proves the existence of chain stable networks and, by using fixed point methods, he shows that the set of chain stable networks is a nonempty lattice. Furthermore, he proves that there exists a consumer optimal network and an initial supplier optimal network, similar to the F -optimal and W -optimal matchings in other models. Finally, for the case in which each agent can be the seller (buyer) in at most one contract, he shows that the set of chain stable networks is equal to the core.

Crawford (2008) proposes to allow offers in the NRMP mechanism to include salaries and demonstrates how the resulting market can generate stable outcomes, which might Pareto dominate the ones in the current form of the NRMP. This model can be seen as an application of the Hatfield and Milgrom (2005) paper.

Another line of investigation that has grown in the last decade concerns the special case of the college admission model in which colleges have fixed preferences, known nowadays as the school choice model. The seminal paper is Sotomayor (1996b) which was motivated by the admission market of economists to graduate centers of economics in Brazil. The students take some tests and each institution places weights on each of these tests in order to rank the students according to the weighted average of the tests. See Ergin and Sönmez (2006), Pathak and Sönmez (2006), and Balinski and Sönmez (1999).

Continuous Two-Sided Matching Model With Additively Separable Utility Functions

The two-sided matching model with *additively separable utility functions* involves two finite and disjoint sets of players which will be denoted by B , with m elements, and Q , with n elements. Each $b \in B$ has a quota $r(b)$ and each $q \in Q$ has a quota $s(q)$ representing the maximum number of partnerships they can form. The main characteristic of this model is that agents are able to negotiate their individual payoffs: If a partnership (b, q) is formed, the partners undertake an activity together that produces a payoff v_{bq} which is divided between them into the payoffs u_{bq} for b and w_{bq} for q , respectively, as a result of a negotiation process. Therefore, an outcome for this game is a matching, along with individual payoffs u_{bq} 's and w_{bq} 's. Dummy players, denoted by 0, are included for technical convenience in both sides of the market. We have that $v_{b0} = v_{0q} = 0$ for all $b \in B$ and $q \in Q$. As for the quotas, a dummy player may form as many partnerships as needed to fill the quotas of the non-dummy players. Then, an allowable set of partners for agent b , with only $r(b) - k$ elements of Q , has k copies of the dummy Q -agent introduced in the model.

A matching μ is feasible if each player is matched to an allowable set of partners. A feasible outcome, denoted by $(u, w; \mu)$, is a feasible matching μ and a pair of payoffs (u, w) , where the individual payoffs of each $b \in B$ and $q \in Q$ are given by the arrays of numbers u_{bq} , with $q \in \mu(b)$, and w_{bq} , with $b \in \mu(q)$, respectively, such that $u_{bq} + w_{bq} = v_{bq}$, $u_{bq} \geq 0$ and $w_{bq} \geq 0$. Consequently, $u_{b0} = u_{0q} = w_{b0} = w_{0q} = 0$ in case these payoffs are defined. We say that the matching μ is compatible with the payoff (u, w) .

The value of μ is $\sum_{q \in Q, b \in \mu(q)} v_{bq}$. The matching μ is *optimal* if it attains the maximum value among all feasible matchings.

This model generates the following game in coalitional function form with side payments. The set of players is $N = B \cup Q$, and the characteristic function v satisfies the following:

1. $v(\emptyset) = 0$.
2. $v(S) = 0$ if $S \subseteq B$ or $S \subseteq Q$.
3. $v(S) \leq v(T)$ if $S \subseteq T$.
4. $v(b, q) = v_{bq}$ for all $(b, q) \in B \times Q$.

For every $b \in B$ and for all sets $S \subseteq Q$ with $|S| \geq r(b)$:

1. $v(b \cup S) = \max \{v(b \cup S'); S' \subseteq S, \text{ and } |S'| = r(b)\}$.
2. (d) implies that for all sets $S \subseteq W$ with $|S| \geq r(b)$, $v(b \cup S) = v(b \cup S')$, and for some $S' \subseteq S$ with $|S'| = r(b)$.

Analogously we define $v(q \cup S)$ for every $q \in Q$ and all sets $S \subseteq B$ with $|S| \geq s(q)$.

The condition that the game has *additively separable utilities* means that for every coalition $S = R \cup T$, $R \subseteq B$, and $T \subseteq Q$:

1. $v(R \cup T) = \max \left\{ \sum_{(b, q) \in R \times T} x_{bq} v_{bq}, \text{ for every feasible assignment } x \right\}$. Consequently, for all $T \subseteq Q$ with $|T| \leq r(b)$ and for all $R \subseteq B$ with $|R| \leq s(q)$, $v(b \cup T) = \sum_{q \in T} v_{bq}$ and $v(q \cup R) = \sum_{b \in R} v_{bq}$.

When the quota of any agent is one, the model is the well-known assignment game introduced in Shapley and Shubik (1972). In this case, the set of individual payoffs of an agent is a singleton, so these payoffs need not be indexed according to the partnerships formed under the matching. Furthermore, the concepts of setwise stability, pairwise stability, and corewise stability are equivalent. *The outcome $(u, w; \mu)$ is stable if it is feasible and $u_b + w_q \geq v_{bq}$ for all pairs (b, q) .*

The existence of stable payoffs for the assignment game was proved in Shapley and Shubik (1972) with the use of linear programming. A different, but also simple proof, is obtained in Sotomayor (2000a) using only combinatorial arguments. Crawford and Knoer (1981) consider a discrete version (as well as the continuous version) of the assignment game and develop a version of the deferred acceptance algorithm of Gale and Shapley to prove the nonemptiness of the set of pairwise-stable payoffs.

The main properties that characterize the set of stable outcomes of the assignment game of Shapley and Shubik are the following:

1. Let (u, w) be some stable payoff. Then μ is an optimal matching if and only if it is compatible with (u, w) .

This result means that the set of stable payoffs is the same under every optimal matching. Then we can concentrate on the payoffs of the agents rather than on the underlying matching.

1. Let (u, w) and (u', w') be stable payoffs. Then $u \geq u'$ if and only if $w' \geq w$.

That is, there exists an opposition of interests between the two sides of the market along the whole set of stable payoffs.

1. If an agent is unmatched under some stable outcome, then he/she gets a zero payoff under any other stable outcome.

This means, for example, that if a worker is unemployed under some stable outcome, then this worker will get a zero salary under any other stable outcome.

1. The set of stable payoffs forms a convex and compact lattice under the partial orders $\geq B$ and $\geq Q$.

The partial order $\geq B$ on the set of stable payoffs is defined as follows: $(u, w) \geq B(u', w')$ if $u_b \geq u'_b$ for all $b \in B$. Property B2 implies that $w_q \leq w'_q$ for all $q \in Q$, so this partial order is well defined. The partial order $\geq Q$ is symmetrically defined. Then, $(u, w) \vee_B(u', w') = (\max\{u, u'\}, \min\{w, w'\})$ and $(u, w) \wedge_B(u', w') = (\min\{u, u'\}, \max\{w, w'\})$.

The lattice property implies that there exist a maximal element and a minimal element in the set of stable payoffs. The fact that the lattice is complete implies the uniqueness of these extreme points. Thus, there exists one and only one stable payoff (u^*, w^*) and one and only one stable payoff (u_*, w_*) such that $(u^*, w^*) \geq B(u, w)$ and

$(u_*, w_*) \geq Q(u, w)$ for all stable payoffs (u, w) . That is:

1. There is a B -optimal stable payoff (u^*, w^*) with the property that for any stable payoff (u, w) , $u^* \geq u$ and $w^* \leq w$; there is a Q -optimal stable payoff (u_*, w_*) with symmetrical properties.
2. The set of stable payoffs equals the core and the set of competitive equilibrium payoffs.

Excluding property B3 which follows from property 1 of Demange and Gale (1985), all the other properties were first proved in Shapley and Shubik (1972).

The general quota case is a version of the model studied in Crawford and Knoer (1981) and was first presented in Sotomayor (1992) in the context of a labor market of firms and workers. Under this approach, the number $r(b)$ is the maximum number of workers firm b can hire, the number $s(q)$ is the maximum number of jobs worker q can take, and the number v_{bq} is the productivity of worker q in firm b . The natural cooperative solution concept is that of setwise stability which is shown to be equivalent to the concept of pairwise stability. Then, the feasible outcome $(u, w; \mu)$ is setwise stable, if $u_b(\min) + w_q(\min) \geq v_{bq}$ for all pairs (b, q) with $q \notin \mu(b)$, where $u_b(\min)$ is the smallest individual payoff of firm b and $w_q(\min)$ is the smallest individual payoff of worker q .

The existence of setwise-stable payoffs for this model was proved in Sotomayor (1992, 1999a) through the use of linear programming.

Another interpretation of this model considers a buyer-seller market: B is a set of buyers and Q is a set of sellers. Buyers are interested in sets of objects owned by different sellers and each seller owns a set of identical objects. The number $r(b)$ is the number of objects buyer b is allowed to acquire, the number $s(q)$ is the number of identical objects seller q owns, and the number v_{bq} is the amount of money buyer b considers to pay for an object of seller q . We say that v_{bq} is the value of object q (object owned by seller q) to buyer b . An artificial null object, 0, owned by the dummy seller, whose value is zero to all buyers and

whose price is always zero, is introduced for technical convenience. Under this approach, a buyer will be assigned to an allowable set of objects at a feasible allocation, meaning that he/she is matched to the set of sellers who own the objects in the given set.

Given a price vector $p \in R_+^\sigma$, with $\sigma \equiv \sum_{q \in Q} s(q)$,

the preferences of buyers over objects are completely described by the numbers v_{bq} 's: For any two allowable sets of objects S and S' , buyer b prefers S to S' at prices p if his/her total payoff when he/she buys S is greater than his/her total payoff when he/she buys S' . He/she is indifferent between these two sets if he/she gets the same total payoff with both sets. Usually, given the prices of the objects, buyers demand their favorite allowable sets of objects at those prices. The set of such allowable sets is called the demand set of buyer b at prices p . *An equilibrium is reached if every buyer is assigned to an allowable set of objects of her demand set, every seller with a positive price sells all of his items, and the number of objects in the market is enough to meet the demand of all buyers.* The solution concept that captures this intuitive idea of equilibrium is that of *competitive equilibrium payoff* defined in Sotomayor (2007a) as an extension of the concept of competitive equilibrium price for the assignment game given in Demange et al. (1986). *Formally, $(u, p; \mu^*)$ is a competitive equilibrium outcome if (i) it is feasible; (ii) μ^* is a feasible allocation such that, if $\mu^*(b) = S$, then S is in the demand set of b at prices p for all $b \in B$; and (iii) $p_q = 0$ if object q is left unsold.*

If $(u, p; \mu^*)$ is a competitive equilibrium outcome, we say that (u, p) is a competitive equilibrium payoff, (p, μ^*) is a competitive equilibrium and p is a *competitive equilibrium price* or an *equilibrium price* for short.

One characteristic of the additively separable utility function is that if a buyer demands a set A of objects at prices p and some of these objects have their prices raised, then the buyer will continue to want to buy the objects in A whose prices were not changed. That is, the function $v(\{b\} \cup A)$ over all allowable sets A of partners for b satisfies the *gross substitute condition*. Kelso and Crawford

(1982) formulated a discrete and a continuous many-to-one matching model where the functions $v(\{b\} \cup A)$ satisfy the gross substitute condition and are not necessarily additively separable. In this model, the core, the set of setwise-stable payoffs, and the set of competitive equilibrium payoffs coincide and are nonempty. These authors prove, through an example, that without this condition, the core may be empty.

A consequence of the competitive equilibrium concept for the many-to-many case with additively separable utility functions is that sellers do not discriminate buyers under a competitive equilibrium payoff, as they might do under a stable outcome. The competitive equilibrium payoffs for this model are characterized as *the setwise-stable payoffs where every seller has identical individual payoffs*. It is interesting to point out that if the identical objects are owned by different sellers, they need not be sold at the same price unless the two sellers have the same number of objects and the selling price is the minimum competitive equilibrium price (Sotomayor 2007a).

A stable (respectively, competitive equilibrium) payoff is called a B -optimal stable (respectively, competitive equilibrium) payoff if every agent in B weakly prefers it to any other stable (respectively, competitive equilibrium) payoff. That is, the B -optimal stable (respectively, competitive equilibrium) payoff gives to each agent in B the maximum total payoff among all stable (respectively, competitive equilibrium) payoffs. Similarly we define a Q -optimal stable (respectively, competitive equilibrium) payoff.

The existence and uniqueness of the B -optimal and of the Q -optimal stable payoffs are proved in Sotomayor (1999a) by showing that the set of stable payoffs is a lattice under two conveniently defined partial orders. This result runs into the difficulties of defining a partial order relation in the set of stable payoffs, due to the fact that, on the one hand, the arrays of individual payoffs are unordered sets of numbers indexed according to the current matching and, on the other hand, the agents' preferences do not define a partial order relation, since they violate the antisymmetric property.

To solve this problem, Sotomayor (1999a) defines a partnership (b, q) to be nonessential if it occurs in some but not all optimal matchings and essential if it occurs in all optimal matchings. Then, two matchings differ only by their nonessential partnerships. According to Theorem 1 of that paper, (i) *in every stable outcome, a player gets the same payoff in any nonessential partnership; furthermore this payoff is less than or equal to any other payoff the player gets under the same outcome*; (ii) *given a stable outcome $(u, w; \mu)$ and a different optimal matching μ' , we can reindex the u_{bq} 's and w_{bq} 's according to μ' and still get a stable outcome*.

Therefore, the array of individual payoffs of a player can be represented by a vector in a Euclidean space whose dimension is the quota of the given player. The first coordinates are the payoffs that the player gets from his essential partners (if any), following some given ordering. The remaining coordinates (if any) are equal to a number which represents the payoff the player gets from all his nonessential partners. This representation is clearly independent of the matching, so any optimal matching is compatible with a stable payoff. Hence, by ordering the players in B (respectively, Q), we can immerse the stable payoffs of these players in a Euclidean space, whose dimension is the sum of the quotas of all players in B (respectively, Q). Then, the natural partial order relation of this Euclidean space induces the partial order relation $\geq B$ (respectively, $\geq Q$) in the set of stable payoffs. We say that $(u, w) \geq B(u', w')$ if the vector of individual payoffs of any buyer, under (u, w) , is greater than or equal to his/her vector of individual payoffs under (u', w') . Similarly we define $(u, w) \geq Q(u', w')$.

The main results of Sotomayor (1999a) are that, under the vectorial representation of the stable payoffs, properties B1, B2, B3, B4, and B5 hold for the general many-to-many case.

An implication of property B2 is the conflict of interests that exists between the two sides of the market with respect to two comparable stable payoffs. That is, if payoffs (u, w) and (u', w') are stable and comparable, then for all $(b, q) \in B \times Q$, we have that b 's total payoff under the first outcome

is greater than b 's total payoff under the second outcome if and only if q 's total payoff under the second outcome is greater than q 's total payoff under the first outcome. From property B3, if a seller has some unsold object under a stable outcome, then one of his/her individual payoffs will be zero under any other such outcome.

Even though the preferences of the players do not define the partial orders $\geq B$ and $\geq Q$, the property stated in B4 is of interest because the two extreme points of the lattice have an important meaning for the model. The extreme points of the lattice are precisely the B -optimal and the Q -optimal stable payoffs. Also every buyer weakly prefers any stable payoff to the Q -optimal stable payoff and any seller weakly prefers any stable payoff to the B -optimal stable payoff.

Indeed, the set of competitive equilibrium payoffs is a sublattice of the set of stable payoffs. This connection is given by the following theorem of Sotomayor (2007a) that states that *the set of competitive equilibrium payoffs is contained in the set of stable payoffs and is a nonempty and complete lattice under the partial order $\geq B$ (respectively, $\geq Q$) whose supremum (respectively, infimum) is B optimal and whose infimum (respectively, supremum) is Q optimal*.

The idea of the proof is that the set of competitive equilibrium payoffs can be obtained by "shrinking" the set of stable payoffs through the application of a convenient isotone (order preserving) map f whose fixed points are exactly the competitive equilibrium payoffs. The desired result is concluded via the *algebraic fixed point theorem* due to Alfred Tarski (1955).

It is also proved in Sotomayor (2007a) that *the B -optimal stable payoff is a fixed point of f , so the B -optimal stable payoff is the B -optimal competitive equilibrium payoff*.

As for property B6, Sotomayor (2003a) shows that the core coincides with the set of pairwise-stable payoffs in the many-to-one case where sellers have a quota of one. Since a seller only owns one object, then he cannot discriminate the buyers, so the core coincides with the set of competitive equilibrium payoffs in this model. Thus, the set of competitive equilibrium payoffs is a lattice in this model. The same result is reached in Gül and Stacchetti (1999)

for the many-to-one case in which the utilities satisfy the gross substitute condition.

However, in the general quota case under additively separable utilities, the core may be bigger than the set of stable payoffs, which in its turn contains and may contain properly the set of competitive equilibrium payoffs, as it is illustrated in the example below from Sotomayor (2007a). This example also shows that the core may not be a lattice and, the polarization of interests, observed in the sets of stable payoffs and of competitive equilibrium payoffs, does not always carry over to the core payoffs: The best core payoff for the buyers is not necessarily the worst core payoff for the sellers.

Example 5 (Sotomayor 2007a) Consider the following situation. The B players will be called firms and the Q players will be called workers. There are two firms, b and b' , and two workers q and q' . Each firm may employ and wants to employ both workers; worker q may take, at most, one job and worker q' may work and wants to work for both firms. The first row of matrix v is (3,2) and the second one is (3,3).

There are two optimal matchings: μ and μ' , where $\mu(b) = \{q, q'\}$, $\mu(b') = \{q', 0\}$ and $\mu'(b) = \{q, q'\}$, $\mu'(b') = \{q', 0\}$. The core is described by the set of individual payoffs (u, w) whose total payoffs (U, W) , satisfy the following system of inequalities: $0 \leq U_b \leq 2$, $0 \leq U_{b'} \leq 3$; $W_q + W_{q'} \geq 3$, $W_{q'} - W_q \leq 2$, $1 \leq W_q \leq 3$. It is not hard to see that the outcome (u, w) is stable if and only if seller q always gets payoff $w_q = 3$ and seller q' gets individual payoffs $w_{bq'} \in [0, 2]$ and $w_{b'q'} \in [0, 3]$; the individual payoffs of buyers b and b' are given by $(u_{bq} = 0, u_{bq'} = 2 - w_{bq'})$ and $(u_{b'q'} = 3 - w_{b'q'}, u_{b'0} = 0)$, respectively.

To see that corewise stability is not adequate to define the cooperative equilibrium for this market, let $(u, w; \mu)$ be such that $u_{bq} = 1$, $u_{bq'} = 1$, $u_{b'q'} = 1$, $u_{b'0} = 0$; $w_{bq} = 2$, $w_{bq'} = 1$, $w_{b'q'} = 2$. That is, firm b hires workers q and q' obtains from each one of them a profit of one and pays two to q and one to q' ; firm b' hires worker q' at a salary of two and obtains a profit of one. Observe that b' has quota of


two, so it has one unfilled position. It happens that b' can pay more than two to q . Thus, if agents can communicate with each other and behave cooperatively, this outcome will not occur, because worker q will not accept to receive only two from firm b , since she knows that she can get more than two by working with firm b' . Hence, this outcome cannot be a cooperative equilibrium. Observe that $2 = u_{b'0} + w_{bq} < v_{b'q} = 3$, so this outcome is not stable. On the other hand, it is in the core. In fact, if there is a blocking coalition, then it must contain $\{b', q\}$. These agents cannot increase their total payoffs by themselves; b' needs to hire both workers. However $\{b', q, q'\}$ does not block the outcome, because q' is worse off by taking only one job. Nevertheless, the coalition of all agents does not block the outcome, since b loses worker q , so it will be worse off.

Now, consider the outcome $(u', w'; \mu)$, where $u_{bq'} = 0$, $u_{bq''} = 1$, $u_{b'q'} = 1$, $u_{b'0'} = 0$; $w_{bq'} = 3$, $w_{bq''} = 1$, and $w_{b'q'} = 2$. Firm b' cannot offer more than three to worker q , so the structure of the outcome cannot be ruptured. Then, although both outcomes $(u, w; \mu)$ and $(u', w'; \mu)$ are corewise stable, only the second one can be expected to occur, so only this outcome is a cooperative equilibrium. Our explanation for this is that only $(u', w'; \mu)$ is stable.

The connection between the core, the set of stable payoffs, and the set of competitive payoffs, exhibited in this example, can be better understood via Fig. 1. Now, if the reader prefers, sets B and Q are better interpreted as being the set of buyers and the set of sellers, respectively.

In Fig. 1, $C(W)$ is the set of seller's total payoffs, which can be derived from any core payoff. The segment OP' is the set of the seller's total payoffs which can be derived from any stable payoff. That is, $(W_q, W_{q'}) \in OP'$ if and only if there is a stable outcome $(u, w; \mu)$ such that $W_q = w_{bq}$ and $W_{q'} = w_{bq'} + w_{b'q'}$. The segment OP is the set of seller's total payoffs, which can be derived from any competitive equilibrium price. That is, $(W_q, W_{q'}) \in OP$ if and only if there is a competitive equilibrium price p such that $W_q = p_q$ and $W_{q'} = p_{q'} + p_{q'}$.

We can see that $C(W)$ is bigger than OP' which, in its turn, is bigger than OP . The point $(2, 3) \in$



Two-Sided Matching Models, Fig. 1 Core, stable, and competitive equilibrium payoffs in Example 5

$C(W) - OP'$. It corresponds to the outcome (u, w, μ) described above that is in the core but is not stable.

It is clear in Fig. 1 that $C(W)$ is not a lattice, so the core is not a lattice under neither $\geq B$ nor $\geq Q$. In fact, the outcome which corresponds to the point $(3,0)$ gives the individual payoff of 3 to q and two individual payoffs of zero to q' . On the other hand, the core outcome that corresponds to $(2,3)$ gives the individual payoff of two to q . Then, the infimum (respectively, supremum) of these two core payoffs under $\geq Q$ (respectively, $\geq B$) gives payoff two to seller q and two individual payoffs of zero to seller q' . This payoff corresponds to the vector of total payoffs $(2,0)$, which is not in $C(W)$.

It is evident that the stable payoff corresponding to point $P' = (3,5)$ is the Q -optimal stable payoff. Seller q receives three and seller q' receives two from b and three from b' . Buyers get zero from the sellers. By applying the function f , we obtain the Q -optimal competitive equilibrium payoff, corresponding to point $P = (3,4)$, where q receives three and q' receives two from both buyers.

Point $O = (3,0)$ corresponds to the outcome (u'', w'', μ) , where $u_{bq}'' = 0$, $u_{bq}''' = 2$, $u_{b'q}''' = 3$, $u_{b'0}'' = 0$; $w_{bq}'' = 3$, $w_{bq}''' = 0$, $w_{b'q}''' = 0$. Payoff (u'', w'') is the B -optimal stable payoff, the B -optimal competitive equilibrium payoff, and the B -optimal core payoff. It is the worst stable

payoff and the worst competitive equilibrium payoff for the sellers. However, it is not the worst core payoff for the sellers since it is the best core payoff for q . Indeed, as it can be observed, there is no minimum core payoff for the sellers.

A fruitful line of investigation has been the design of mechanisms to produce a competitive equilibrium price. Demange et al. (1986) propose a generalization of the English auction for the assignment game, which yields the minimum competitive equilibrium price in a finite number of steps. In the same spirit, Sotomayor (2002) presents a descending bid auction mechanism which leads the maximum competitive equilibrium price as a generalization of the Dutch auction. Gül and Stacchetti (2000) obtain the minimum competitive equilibrium price through a generalization of the auction of Demange et al. (1986) to the many-to-one case in which the utility functions satisfy the gross substitute condition. Sotomayor (2006) obtains the same result by considering a dynamic mechanism for the many-to-many case with additively separable utility functions. This mechanism leads to the B -optimal stable payoff. Using the symmetry of the model, the Q -optimal stable payoff is obtained by reverting the roles between buyers and sellers in the mechanism.

A number of works related to the assignment game can be found in the literature. Demange and Gale (1985) generalize the assignment game by allowing agents' preferences to be represented by any continuous utility function in the money variable. For that model, Roth and Sotomayor (1988) generalize a previous result of Rochford (1984), using Tarski's fixed point theorem, and show that the set of fixed points of a "rebargaining" function is a subset of the core, which maintains the lattice structure of the core. Another approach of the many-to-many assignment game with additively separable utilities is treated in Sotomayor (1992). There, agents are not allowed to negotiate their individual payments and act in blocks. An outcome only specifies their total payoffs. In this model, the core coincides with the set of stable payoffs and is not a lattice. Also, pairwise stability is not equivalent to corewise stability. Sotomayor (2003a) considers an extension of the assignment

game to a many-to-many matching model in the context of firms and workers in which the quotas of the agents are not the number of partnerships they are allowed to form. Instead, they are given by the units of labor time they can supply or employ. Bikhchandani and Mamer (1997) analyze the existence of market clearing prices in an exchange economy in which agents have interdependent values over several indivisible objects. Although an agent can be both a buyer and a seller, such an exchange economy can be transformed into a many-to-one matching market where each seller owns only one object and buyers want to buy a bundle of objects, and can be viewed as an extension of the assignment game. See also Demange (1982), Leonard (1983), Perez-Castrillo and Sotomayor (2002), Sotomayor (2003c), Thompson (1980) and Kaneko (1982).

Hybrid One-to-One Matching Model

The hybrid one-to-one matching model is the name given in the literature to a unified model due to Eriksson and Karlander (2000) and inspired in the unification proposed in Roth and Sotomayor (1996). Agents from the marriage market and the assignment market are put together so that they can trade with each other in the same market. We can interpret the hybrid model as being a labor market of firms and workers: P is the set of firms and Q is the set of workers. There are two classes of agents in each set: rigid agents and flexible agents. For each pair $(p, q) \in P \times Q$, there is a number c_{pq} representing the productivity of the pair. If a firm p hires worker q and both agents are flexible, then the number c_{pq} is allocated into salary v_q for the worker and profit u_p for the firm as a result of a negotiation process. If one of the agents is rigid, then the payoffs of the agents are preset and fixed and are part of the job description. In this case, the profit of p and the salary of q will be a_{pq} and b_{pq} , respectively. The definitions of feasible outcome, corewise-stable outcomes, and setwise-stable outcomes are straightforward extensions from the respective concepts for the non-hybrid models.

Therefore, the concepts of setwise stability and corewise stability are equivalent.

This model is motivated by the fact that, in practice, a wide range of real-world matching markets are neither completely discrete nor completely continuous. In the United States, for example, new law school graduates may enter the market for associate positions in private law firms, which negotiate salaries, or they may seek employment as law clerks to federal circuit court judges, which are civil service positions with predetermined fixed salaries. In the market for academic positions and professors, for example, the American universities compete with each other in terms of salaries, while the French public universities offer a preset and fixed salary. In Brazil, new professors may enter the market for permanent positions (with preset and fixed salaries) in federal universities, or they may seek employment in private universities, which do not offer such positions but compensate the entrants with better and negotiable salaries.

Eriksson and Karlander (2000) present an algorithm to find a stable outcome under the assumption that the numbers a_{pq} , b_{pq} , and c_{pq} are integer numbers of some unit. A nonconstructive proof of the existence result without imposing any restriction is provided in Sotomayor (2000a). In this paper, it is proved that, under the assumption that the core, C , is equal to the strong core, C^* , the main properties that characterize the stable payoffs of the marriage and of the assignment models carry over to the hybrid model. That is, for the hybrid model:

1. Let (u, w) and (u', w') be stable payoffs for the hybrid model. If $C = C^*$, then $u \geq u'$ if and only if $w' \geq w$.
2. If $C = C^*$ and an agent is unmatched under some stable outcome then he/she gets a zero payoff under any other stable outcome.
3. If $C = C^*$, then the set of stable payoffs forms a complete lattice under the partial orders $\geq_P$ and $\geq_Q$.
4. If $C = C^*$ then there is a P -optimal stable payoff (u^*, w^*) with the property that for any stable payoff (u, w) , $u^* \geq u$ and $w^* \leq w$; there is a Q -optimal stable payoff (u_*, w_*) with symmetrical properties.

Sotomayor (2007b) studies the hybrid model without imposing the assumption that the core is equal to the strong core. Instead she assumes that the preferences of the rigid agents, as well as the preferences of the flexible agents over rigid agents, are strict. It is shown there that the core of the hybrid model has a nonstandard algebraic structure given by the disjoint union of complete lattices endowed with the properties above. The extreme points of the lattices of the core partition are called *quasi-optimal stable payoffs for firms* and *quasi-optimal stable payoffs for workers*. When the workers are always flexible, then the marriage market is obtained when the flexible firms leave the hybrid market and the assignment game is obtained when the rigid firms leave the hybrid market.

Each subset of the core partition of the hybrid model is obtained as follows. For any matching μ which is compatible with some stable payoff, decompose the market participants into two disjoint subsets. One subset contains all rigid firms and their mates at μ and the other one contains all flexible firms, their mates at μ , and the unmatched workers. Now, fix such a partition of the agents. The desired subset $C(\mu)$ of the core partition is formed with the core payoffs $(u, w; \mu')$ such that all agents in the first set are matched among themselves under μ' and all agents in the second set are matched among themselves under μ' .

Clearly, as the rigid firms exit the hybrid market, the core partition for the corresponding assignment market is reduced to only one set, since any stable matching is compatible with any core payoff by property B1. An analogous result holds as the flexible firms leave the hybrid market, due to the fact that the matched agents in the marriage market are the same at every stable matching, which is implied by property A2. Therefore, as all flexible firms leave the hybrid market or as all rigid firms leave the hybrid market, the restriction of the algebraic structure to the core of the resulting non-hybrid market is that of a complete lattice. Then, the extreme points of the resulting lattice are exactly the firm-optimal and the worker-optimal stable payoffs.

This algebraic structure was used in Sotomayor (2007b) to investigate the comparative effects on the quasi-optimal stable payoffs for firms and on the

quasi-optimal stable payoffs for workers caused by the entrance of rigid firms into the assignment market or by the entrance of flexible firms into the marriage model. The results of that paper can be summarized as follows: *Whether agents are allocated according to a quasi-optimal stable payoff for firms or according to a quasi-optimal stable payoff for workers, it will always be the case that if flexible firms enter the rigid market, no rigid firm will be made better off and no worker will be made worse off; if rigid firms enter the flexible market, no flexible firm will be made better off and no worker will be made worse off.*

Comparative static results of adding agents from the same side to the marriage market or to the assignment market have been obtained in the literature under the assumption that the agents are allocated according to the optimal stable outcome for firms or according to the optimal stable outcome for workers. However, in the approach considered in Sotomayor (2007b):

1. *The firms that are added are different from the firms, which are already in the market.* For example, in the marriage market, where utility is non-transferable, the comparative static adds firms with flexible wages who can transfer utility.
2. *The points which are compared belong to cores with quite distinct algebraic structures.*
3. *There may exist several quasi-optimal stable outcomes for firms and several quasi-optimal stable outcomes for workers in the hybrid model. Despite the multiplicity of these outcomes, all of them reveal the same kind of comparative static effects.*

Therefore, the result above has no parallel in the non-hybrid models.

It is argued in Sotomayor (2007b) that if the resulting core partition is not reduced to only one set when, say, the flexible firms leave the hybrid model, the comparative statics may be meaningless. This happens, for example, if we define a set of the core partition as the set of all stable payoffs compatible with some given matching. Then each lattice of the core partition for the marriage market has only one stable matching, which is both the

supremum and the infimum of the lattice. Of course, the distinctions between, say, the best stable payoff for workers of some lattice of the core partition of the hybrid market and an arbitrary core point of the marriage model cannot be attributed to the entrance of the flexible firms into the marriage market.

Results of comparative statics were originally obtained by Gale and Sotomayor (1983) for the marriage model and the college admission model: *If agents from the same side of the market enter the market, then no agent from this side is better off and no agent of the opposite side is worse off, if any of the two optimal stable matchings prevails.* A similar result was proved by Demange and Gale (1985) for a continuous one-to-one matching model that includes the assignment game.

For the assignment game, Shapley (1962) showed that the optimal stable payoff for an agent weakly decreases when another agent is added to the same side and weakly increases when another agent is added to the other side. Still with regard to the assignment game, Mo (1988) showed that if the incoming worker is allocated to some firm in some stable outcome for the new market, there is a set of agents such that every firm is better off and every worker is worse off in the new market than in the previous one. A symmetric result holds when the incoming agent is a firm. An analogous result is demonstrated by Roth and Sotomayor (1990) for the marriage market.

For the many-to-one matching markets with substitutable preferences, Kelso and Crawford (1982) showed that, within the context of (flexible) firms and workers, *the addition of one or more firms to the market weakly improves the workers' payoffs and the addition of one or more workers weakly improves the firms' payoffs, under the firm-optimal stable allocation.* Similar conclusions were obtained by Crawford (1991) for a many-to-many matching model with strict and substitutable preferences, by comparing pairwise-stable outcomes instead of setwise-stable outcomes.

Incentives

The strategic questions that emerge when a stable revelation mechanism is adopted concerns its

non-manipulability and, for the games induced by the mechanism, the existence of strategic equilibria and the implementability of the set of stable matchings via such equilibria. For the marriage model with strict, the equilibrium analysis of a game induced by a stable matching mechanism leads to the following results:

1. (Impossibility theorem) (Roth and Sotomayor, 1990) *When any stable mechanism is applied to a marriage market in which preferences are strict and there is more than one stable matching, then at least one agent can profitably misrepresent his/her preferences, assuming the others tell the truth (this agent can misrepresent in such a way as to be matched to his/her most preferred achievable mate under the true preferences at every stable matching under the mis-stated preferences).*
2. (Limits on successful manipulation) (Demange et al. 1987). *Let P be the true preferences (not necessarily strict) of the agents, and let P' differ from P in that some coalition C of men and women misstate their preferences. Then there is no matching μ , stable for P' , which is preferred to every stable matching under the true preferences P by all members of C .*

A corollary of this result is due to Dubins and Freedman (1981) which states that *the man-optimal stable matching mechanism is non-manipulable, individually and collectively, by the men:*

1. (Gale and Sotomayor 1985) *When all preferences are strict, let μ be any stable matching for (F, W, P) . Suppose each woman w in $\mu(F)$ chooses the strategy of listing only $\mu(w)$ on her stated preference list of acceptable men (and each man states his true preferences). This is a Nash equilibrium in the game induced by the man-optimal stable matching mechanism (and μ is the matching that results).*
2. (Roth 1984a) *Suppose each man chooses his dominant strategy and states his true preferences and the women choose any set of strategies (preference lists) $P'(w)$ that form a Nash equilibrium for the revelation game induced by the man-optimal stable mechanism. Then the*

- corresponding man-optimal stable matching for (F, W, P') is one of the stable matchings for (F, W, P) .
3. (Gale and Sotomayor 1985) Suppose each man chooses his dominant strategy and states his true preferences and the women truncate their true preferences at the mate they get under the woman-optimal stable mechanism. This profile of preferences is a strong equilibrium for the women in the game induced by the man-optimal stable mechanism (and the woman-optimal stable matching under the true preferences is the matching that results).

Results (3) and (4) imply that the man-optimal stable mechanism implements the core correspondence via Nash equilibria.

For the college admission model with responsive and strict preferences, the theorem of Dubins and Freedman implies that the student-optimal stable mechanism is non-manipulable individually and collectively by the students. Roth (1985b) shows through an example that the college-optimal stable mechanism is manipulable by the colleges due to the fact that the colleges may have a quota greater than one.

Sotomayor (1998, 2000, 2007c) analyzes the strategic behavior of the students in a school choice model where participants have strict preferences over individuals. This paper proves that the college-optimal stable mechanism implements the set of stable matchings via the Nash equilibrium concept. When some other stable mechanism is used, an example shows that the strategic behavior of the students may lead to unstable matchings under the true preferences. A sufficient condition for the stability of the Nash equilibrium outcome is then proved to be that the set of stable matchings for the Nash equilibrium profile is a singleton. A random stable matching mechanism is proposed and the Nash equilibrium concept ex ante is shown to be equivalent to the Nash equilibrium concept ex post of the game induced by such a mechanism. This refinement of the Nash equilibrium concept is called *Nash equilibrium in the strong sense*. Under this equilibrium concept, any stable matching mechanism (and in particular the random stable matching mechanism) implements the set of stable matchings. Also, if the students only play truncations of the true

preferences, any stable matching mechanism implements the student-optimal stable matching via strong equilibrium in the strong sense and Nash equilibrium in the strong sense.

Ergin and Sönmez (2006) and Pathak and Sönmez (2006) analyze the Boston mechanism, which is used to assign students to schools in many cities in the United States and show that students' parents do not have incentives to report preferences truthfully.

Ma (2002) analyzes the strategic behavior of both students and colleges in the college admission model with responsive preferences. This author proves that the set of stable matchings is implemented by any stable mechanism via rematching proof equilibrium and strong equilibrium in truncation strategies at the match point.

The implementability of the set of stable matchings through stable and non-necessarily stable mechanisms has also been investigated by several authors. Alcalde (1996) presents a mechanism for the marriage market closely related to the algorithm of Gale and Shapley, which implements the core correspondence in undominated equilibria. Kara and Sönmez (1996) analyze the problem of implementation in the college admission market. They show that the set of stable matchings is implementable in Nash equilibrium. Nevertheless, no subset of the core is Nash implementable. Romero-Medina (1998) studies the mechanism employed by the Spanish universities to distribute the students to colleges, which can produce unstable matchings for the stated preferences. However, when students play in equilibrium, only stable allocations are reached. Sotomayor (2003b) investigates a mechanism for the marriage model which is not designed for producing stable matchings. Here also the equilibrium outcomes are stable matchings under the true preferences.

For the discrete many-to-one matching model with responsive preferences, Alcalde and Romero-Medina (2000) analyze the following mechanism: Firms announce a set of workers they want to hire. Then each worker selects the firm she wants to work for. This paper proves that such a mechanism implements the set of stable matchings in subgame perfect Nash equilibrium. For the many-to-many case, Sotomayor (2004b) shows that this result does not carry over. This paper proves that

subgame perfect Nash equilibria always exist, while strong equilibria may not exist. The subgame perfect Nash equilibrium outcomes are precisely the pairwise-stable matchings, which may be out of the core when the preferences of the agents in one of the sides are not maximin. Under this condition, the equilibrium outcomes are the setwise-stable matchings and every subgame perfect Nash equilibrium is a strong equilibrium.

By assuming non-strict preferences, Abdulkadiroğlu et al. (2006) show that no mechanism (stable or not and Pareto optimal or not), which is better for the students than the student proposing deferred acceptance algorithm with tie breaking, can be strategy proof.

Sönmez (1999) analyzes a model which he calls the *generalized indivisible allocation problem* that includes the roommate and the marriage markets. He looks for conditions which explain the differences on strategy-proofness results that have been generated in the literature. He shows how some of the results in the literature can be seen as corollaries of his results.

Ehlers and Massó (2004) study Bayesian Nash equilibria of stable mechanisms (such as the NRMP) in matching markets under incomplete information. They show that truth-telling is an equilibrium of the Bayesian revelation game induced by a common belief and a stable mechanism if and only if all the profiles in the support of the common belief have singleton cores.

For the continuous matching models, the idea is to use competitive equilibrium as an allocation mechanism to produce outcomes with the desirable properties of fairness and efficiency. It involves having agents specify their supply and demand functions. The competitive equilibria are then calculated and allocations are made accordingly. Demange (1982) and Leonard (1983) considered the assignment game of Shapley and Shubik and, independently, proved that the allocation mechanism that yields the minimum competitive equilibrium prices is individually non-manipulable by the buyers. Demange and Gale (1985) consider a one-to-one matching model in which the utilities are continuous functions in the money variable and not necessarily additively separable. These authors prove a sort of non-manipulability theorem which states that *if the mechanism which produces the*

buyer-optimal stable payoff is adopted, then no coalition of buyers by falsifying demands can achieve, only through the mechanism, higher payoffs to all of its members. We added *only through the mechanism* because this model allows monetary transfers within any coalition. As in the marriage model, this result is an immediate consequence of a more general theorem due to Sotomayor (1986) which states the following: *Let $(u', w'; \mu)$ be any stable outcome for the market M' where $B' \cup Q'$ is the set of agents who misrepresent their utility functions. Let (u^*, w^*) be the true payoff under $(u', w'; \mu)$. Then, there exists a stable payoff (u, w) for the original market such that $u_b \geq u^* \cdot b$ for at least one b in B' or $w_q \geq w'_q$ for at least one q in Q' .*

Demange and Gale (1985) also addresses the strategic behavior by the sellers when the mechanism produces the buyer-optimal stable payoff. These authors show that by specifying their supply functions appropriately, the sellers can force, by strong Nash equilibrium strategies, the payoff to be given by the maximum rather than the minimum equilibrium price. Under the assumption that the sellers only manipulate their reservation prices then, if a profile of strategies does not give the maximum equilibrium price allocation, then either some seller is using a dominated strategy or the strategy profile is not a Nash equilibrium. For this model, Sotomayor (1986) proves that the outcome produced by a Nash equilibrium strategy is stable for the original market.

Sotomayor (2004a) considers, for the assignment game of Shapley and Shubik, the strategic games induced by a class of market clearing price mechanisms. In these procedures, buyers and sellers in different stages reveal their demand and supply functions and a competitive equilibrium is produced by the mechanism. For each vector of reservation prices selected by the sellers, the buyers play the subgame that starts and can force the buyer-optimal stable payoff through Nash equilibrium strategies. However sellers can reverse this outcome by forcing the subgame perfect equilibrium allocation to be the seller-optimal stable payoff for the original market.

Kamecke (1989) and Perez-Castrillo and Sotomayor (2002) consider the assignment game of Shapley and Shubik. The former paper presents two mechanisms for this market. In the first one,

agents act simultaneously. In the second game, the strategies are chosen sequentially. These mechanisms implement the social choice correspondences that yield the core and the optimal stable payoff for the sellers, respectively. The second paper analyzes a sequential mechanism, which implements the social choice correspondence that yields the optimal stable payoff for the sellers.

Future Directions

In this section, we present some directions for future investigations and some open problems which have intrigued matching theorists.

1. The discrete two-sided matching models with non-necessarily strict preferences have been explored very little in the literature. In the discrete models under strict preferences and in the continuous models, due to the fact that there is no weak blocking pairs, the set of Pareto-stable outcomes coincides with the set of setwise-stable outcomes. However, under weak preferences, setwise-stable matchings may not be Pareto optimal. Sotomayor (2008) proposes that in this case the Pareto-stability concept, which requires that the matching is stable and Pareto optimal, should be considered the natural solution concept. The justification for the Pareto stability concept relies in the argument that in a decentralized setting, where agents freely get together in groups, recontracts between pairs of agents already allocated according to a stable matching leading to a (weak) Pareto improvement of the original matching should be allowed. Thus, weak blockings can upset a matching once they come from the grand coalition.

We think that the study of the discrete two-sided matching models with non-necessarily strict preferences and the search for algorithms to produce the Pareto-stable matchings is a new and interesting line of investigation.

2. Consider the hybrid model where no worker is rigid. If rigid firms enter the flexible market or flexible firms enter the rigid market, then no firm gains and no worker loses if a quasi-optimal stable outcome for one of the sides always prevails. Suppose now that some rigid firm becomes flexible or some flexible firm

becomes rigid. What kind of comparative static effect is caused by this change in the market?

3. One line of investigation not yet explored in the literature concerns the incentives faced by the agents in the hybrid model when some stable allocation mechanism is used.
4. Consider the discrete many-to-many matching market with substitutable preferences where a matching μ is feasible if $Ch_y(\mu(y)) = \mu(y)$ for every agent y . Is the core always nonempty for this model?
5. Consider the assignment game of Shapley and Shubik in the context of buyers and sellers. Consider a sealed bid auction in which the buyers select a monetary value for each of the items. The auctioneer then chooses a competitive equilibrium price vector for the profile of selected values, according to some preset probability distribution. It is of theoretical interest the investigation of the buyers' strategic behavior.
6. We know that the core of the many-to-many assignment model of Sotomayor (1992) in which the agents negotiate in blocks is not a lattice. However, a problem that is still open is to know if the optimal stable payoffs for each side of the market always exist.

Bibliography

- Abdulkadiroglu A, Sönmez T (2003) School choice: a mechanism design approach. *Am Econ Rev* 93(3):729–747
- Abdulkadiroglu A, Pathak P, Roth A, Sönmez T (2006) Changing the Boston school choice mechanism: strategy-proofness as equal access. Working paper. Boston College and Harvard University, Boston
- Abeledo H, Isaak G (1991) A characterization of graphs which assure the existence of stable matchings. *Math Soc Sci* 22(1):93–96
- Adachi H (2000) On a characterization of stable matchings. *Econom Lett* 68(1):43–49
- Alcalde J (1996) Implementation of stable solutions to marriage problems. *J Econom Theory* 69(1):240–254
- Alcalde J, Romero-Medina A (2000) Simple mechanisms to implement the core of college admissions problems. *Games Econom Behav* 31(2):294–302
- Alcalde J, Pérez-Castrillo D, Romero-Medina A (1998) Hiring procedures to implement stable allocations. *J Econom Theory* 82(2):469–480
- Balinski M, Sönmez T (1999) A tale of two mechanisms: student placement. *J Econom Theory* 84(1):73–94

- Bardella F, Sotomayor M (2006) Redesign and analysis of an admission market to the graduate centers of economics in Brazil: a natural experiment in market organization, working paper. Universidade de São Paulo, São Paulo
- Bikhchandani S, Mamer J (1997) Competitive equilibrium in an exchange economy with indivisibilities. *J Econom Theory* 74(2):385–413
- Birkhoff G (1973) Lattice theory, vol v. 25. Colloquium publications, American Mathematical Society, Providence
- Blair C (1988) The lattice structure of the Set of stable matchings with multiple partners. *Math Oper Res* 13(4):619–628
- Chung K (2000) On the existence of stable roommate matchings. *Games Econom Behav* 33(2):206–230
- Crawford V (1991) Comparative statics in matching markets. *J Econom Theory* 54(2):389–400
- Crawford V (2008) The flexible-salary match: a proposal to increase the salary flexibility of the national resident matching program. *J Econom Behav Organ* 66(2):149–160
- Crawford V, Knoer E (1981) Job matching with heterogeneous firms and workers. *Econometrica* 49(2):437–450
- Demange G (1982) Strategyproofness in the assignment market game, working paper. Ecole Polytechnique, Laboratoire D'Econometrie, Paris
- Demange G, Gale D (1985) The strategy structure of two-sided matching markets. *Econometrica* 53(4):873–888
- Demange G, Gale D, Sotomayor M (1986) Multi-item auctions. *J Political Econom* 94(4):863–872
- Demange G, Gale D, Sotomayor M (1987) A further note on the stable matching problem. *Discrete Appl Math* 16(3):217–222
- Dubins L, Freedman D (1981) Machiavelli and the Gale-Shapley algorithm. *Am Math Mon* 88(7):485–494
- Dutta B, Massó J (1997) Stability of matchings when individuals have preferences over colleagues. *J Econom Theory* 75(2):464–475
- Echenique F, Oviedo J (2004) Core many-to-one matchings by fixed-point methods. *J Econom Theory* 115(2):358–376
- Echenique F, Oviedo J (2006) A theory of stability in many-to-many matching markets. *Theor Econ* 1(2):233–273
- Echenique F, Yenmez M (2007) A solution to matching with preferences over colleagues. *Games Econom Behav* 59(1):46–71
- Eeckhout J (2000) On the uniqueness of stable marriage matchings. *Econom Lett* 69(1):1–8
- Ehlers L, Massó J (2004) Incomplete information and small cores in matching markets, working paper. CREA, Barcelona
- Ergin H, Sönmez T (2006) Games of school choice under the Boston mechanism. *J Public Econom* 90(1–2):215–237
- Eriksson K, Karlander J (2000) Stable matching in a common generalization of the marriage and assignment models. *Discrete Math* 217(1):135–156
- Gale D, Shapley L (1962) College admissions and the stability of marriage. *Am Math Mon* 69(1):9–15
- Gale D, Sotomayor M (1983, 1985) Some remarks on the stable matching problem. *Discret Appl Math* 11:223–232
- Gale D, Sotomayor M (1985) Ms. Machiavelli and the stable matching problem. *Am Math Mon* 92(4):261–268
- Gül F, Stacchetti E (1999) Walrasian equilibrium with gross substitutes. *J Econom Theory* 87(1):95–124
- Gül F, Stacchetti E (2000) The english auction with differentiated commodities. *J Econom Theory* 92(1):66–95
- Gusfield D (1988) The structure of the stable roommate problem: efficient representation and enumeration of all stable assignments. *SIAM J Comput* 17:742–769
- Hatfield J, Milgrom P (2005) Matching with contracts. *Am Econ Rev* 95(4):913–935
- Irving R (1985) An efficient algorithm for the stable roommates problem. *J Algorithms* 6:577–595
- Kamecke U (1989) Non-cooperative matching games. *Int J Game Theory* 18(4):423–431
- Kaneko M (1982) The central assignment game and the assignment markets. *J Math Econom* 10(2–3):205–232
- Kara T, Sönmez T (1996) Nash implementation of matching rules. *J Econom Theory* 68(2):425–439
- Kara T, Sönmez T (1997) Implementation of college admission rules. *J Econom Theory* 9(2):197–218
- Kelso A Jr, Crawford V (1982) Job matching, coalition formation, and gross substitutes. *Econometrica* 50(6):1483–1504
- Kesten O (2004) Student placement to public schools in the US: two new solutions, working paper. University of Rochester, Rochester
- Kesten O (2006) On two competing mechanisms for priority-based allocation problems. *J Econom Theory* 127(1):155–171
- Knuth D (1976) Marriageables. Les Presses de l'Université de Montréal, Montréal
- Konishi H, Ünver MU (2006) Credible group stability in many-to-many matching problems. *J Econom Theory* 127(1):57–80
- Kraft C, Pratt J, Seidenberg A (1959) Intuitive probability on finite sets. *Ann Math Stat* 30(2):408–419
- Leonard H (1983) Elicitation of honest preferences for the assignment of individuals to positions. *J Political Econ* 91(3):461–479
- Ma J (2002) Stable matchings and the small core in Nash equilibrium in the college admissions problem. *Rev Econ Des* 7(2):117–134
- Martínez R, Massó J, Neme A, Oviedo J (2001) On the lattice structure of the set of stable matchings for a many-to-one model. *Optimization* 50(5):439–457
- Martínez R, Massó J, Neme A, Oviedo J (2004) An algorithm to compute the full set of many-to-many stable matchings. *Math Soc Sci* 47(2):187–210
- McVitie D, Wilson L (1970) Stable marriage assignment for unequal sets. *BIT Numer Math* 10(3):295–309
- Mo J (1988) Entry and structures of interest groups in assignment games. *J Econom Theory* 46(1):66–96
- Ostrovsky M (2008) Stability in supply chain networks. *Am Econ Rev* 98(3):897–923
- Pathak P, Sönmez T (2006) Leveling the playing field: sincere and strategic players in the Boston mechanism, working paper. Boston College and Harvard University, Boston

- Pérez-Castrillo D, Sotomayor M (2002) A simple selling and buying procedure. *J Econom Theory* 103(2):461–474
- Pycia M (2007) Many-to-one matching with complementarities and peer effects, working paper. Penn State working paper, Pennsylvania
- Rochford S (1984) Symmetrically pairwise-bargained allocations in an assignment market. *J Econom Theory* 34(2):262–281
- Romero-Medina A (1998) Implementation of stable solutions in a restricted matching market. *Rev Econ Des* 3(2):137–147
- Roth A (1982) The economics of matching: stability and incentives. *Math Oper Res* 7(4):617–628
- Roth A (1984a) Misrepresentation and stability in the marriage problem. *J Econom Theory* 34(2):383–387
- Roth A (1984b) Stability and polarization of interests in job matching. *Econometrica* 52(1):47–58
- Roth A (1984c) The evolution of the labor market for medical interns and residents: a case study in game theory. *J Political Econ* 92(6):991–1016
- Roth A (1985a) Conflict and coincidence of interest in job matching: some new results and open questions. *Math Op Res* 10(3):379–389
- Roth A (1985b) The college admissions problem is not equivalent to the marriage problem. *J Econom Theory* 36(2):277–288
- Roth A (1986) On the allocation of residents to rural hospitals: a general property of two-sided matching markets. *Econometrica* 54(2):425–427
- Roth A, Sotomayor M (1988) Interior points in the core of two-sided matching markets. *J Econom Theory* 45(1):85–101
- Roth A, Sotomayor M (1989) The college admissions problem revisited. *Econometrica* 57(3):559–570
- Roth A, Sotomayor M (1990) Two-sided matching: a study in game-theoretic modeling and analysis. In: *Econometric society monographs*, vol 18. Cambridge University Press, New York
- Roth A, Sotomayor M (1996) Stable outcomes in discrete and continuous models of two-sided matching: a unified treatment. *Braz Rev Econom* 16:1–4
- Shapley L (1962) Complements and substitutes in the optimal assignment problem. *Navals Res Logist Q* 9:45–48
- Shapley L, Shubik M (1972) The assignment game I: the core. *Int J Game Theory* 1(1):111–130
- Sönmez T (1999) Strategy-proofness and essentially single-valued cores. *Econometrica* 67(3):677–689
- Sotomayor M (1986) On incentives in a two-sided matching market, working paper. Depart of Mathematics, PUC/RJ, Rio de Janeiro
- Sotomayor M (1992) The multiple partners game in equilibrium and dynamics. In: Majumdar M (ed) *Essays in honour of David Gale*. Macmillan, New York, pp 322–336
- Sotomayor M (1996a) A non-constructive elementary proof of the existence of stable marriages. *Games Econom Behav* 13(1):135–137
- Sotomayor M (1996b) Admission mechanisms of students to colleges. A game-theoretic modeling and analysis. *Braz Rev Econom* 16(1):25–63
- Sotomayor M (1998) The strategy structure of the college admissions stable mechanisms. In: *Annals of Jornadas Latino Americanas de Teoria Econômica*, San Luis, Argentina, 2000; *First World Congress of the Game Theory Society*, Bilbao, 2000; *4th Spanish Meeting*, Valencia, 2000; *International Symposium of Mathematical Programming*, Atlanta, 2000; *World Congress of the Econometric Society*, Seattle, 2000, http://www.econ.fea.usp.br/marilda/artigos/roommates_1.doc
- Sotomayor M (1999a) The lattice structure of the set of stable outcomes of the multiple partners assignment game. *Int J Game Theory* 28(4):567–583
- Sotomayor M (1999b) Three remarks on the many-to-many stable matching problem. *Math Soc Sci* 38(1):55–70
- Sotomayor M (2000a) Existence of stable outcomes and the lattice property for a unified matching market. *Math Soc Sci* 39(2):119–132
- Sotomayor M (2000) Reaching the core through college admissions stable mechanisms. In: *Annals of abstracts of the following congresses: International Conference on Game Theory*, 2001, Stony Brook; *Annals of Brazilian Meeting of Econometrics*, Salvador, Brazil, 2001; *Latin American Meeting of the Econometric Society*, Buenos Aires, Argentina, 2001, http://www.econ.fea.usp.br/marilda/artigos/reaching_%20core_random_stable_allocation_mechanisms.pdf
- Sotomayor M (2002) A simultaneous descending bid auction for multiple items and unitary demand. *Revista Brasileira Economia* 56:497–510
- Sotomayor M (2003a) A labor market with heterogeneous firms and workers. *Int J Game Theory* 31(2):269–283
- Sotomayor M (2003b) Reaching the core of the marriage market through a non-revelation matching mechanism. *Int J Game Theory* 32(2):241–251
- Sotomayor M (2003c) Some further remark on the core structure of the assignment game. *Math Soc Sci* 46:261–265
- Sotomayor M (2004a) Buying and selling strategies in the assignment game, working paper. Universidade São Paulo, São Paulo
- Sotomayor M (2004b) Implementation in the many-to-many matching market. *Games Econom Behav* 46(1):199–212
- Sotomayor M (2005) The roommate problem revisited, working paper. Universidade São Paulo, São Paulo. http://www.econ.fea.usp.br/marilda/artigos/THE_ROOMMATE_PROBLEM_REVISITED_2007.pdf
- Sotomayor M (2006) Adjusting prices in the many-to-many assignment game to yield the smallest competitive equilibrium price vector, working paper. Universidade de São Paulo, São Paulo. http://www.econ.fea.usp.br/marilda/artigos/A_dynam_stable-mechan_proof_lemma_1.pdf
- Sotomayor M (2007a) Connecting the cooperative and competitive structures of the multiple partners assignment game. *J Econom Theory* 134(1):155–174
- Sotomayor M (2007b) Core structure and comparative statics in a hybrid matching market. *Games Econom Behav* 60(2):357–380

- Sotomayor M (2007c) The stability of the equilibrium outcomes in the admission games induced by stable matching rules. *Int J Game Theory* 36(3–4):621–640
- Sotomayor M (2008) The pareto-stability concept is a natural solution concept for the discrete matching markets with indifferences, working paper. Universidade São Paulo, São Paulo. http://www.econ.fea.usp.br/marilda/artigos/ROLE_PLAYED_SIMPLE_OUTCOMES_STABLE_COALITI_3.pdf
- Tan J (1991) A necessary and sufficient condition for the existence of a complete stable matching. *J Algorithms* 12(1):154–178
- Tarski A (1955) A lattice-theoretical fixpoint theorem and its applications. *Pacific J Math* 5(2):285–309
- Thompson G (1980) Computing the core of a market game. In: Fiacco AA, Kortane K (eds) *Extremal methods and systems analysis*, vol 174. Springer, New York, pp 312–324

Market Design

Fuhito Kojima, Fanqi Shi and Akhil Vohra
Department of Economics, Stanford University,
Stanford, CA, USA

Article Outline

Introduction
Two-Sided Matching
One-Sided Matching
Applications
Conclusion
Bibliography

Introduction

Matching theory in economics began with the seminal contribution by Gale and Shapley (1962). Ever since, the theory has advanced considerably and has been applied to an increasing number of economic problems. Notably, it has proved useful in helping designs of mechanisms in a variety of markets. Examples include medical match (Roth 1984; Roth and Peranson 1999) and other entry-level labor markets (Roth 1991), school choice (Abdulkadiroğlu and Sönmez 2003), course allocation in education (Sönmez and Ünver 2010; Budish and Cantillon 2012), and organ donation (Roth et al. 2004, 2005, 2007). Application of matching theory to these and other practical problems is known as “market design.” Although market design is often used to refer to other types of research as well, in this article, we focus on market design as application of matching theory.

This paper describes matching theory and its applications. We begin by describing standard models in two-sided and one-sided (object allocation) models in some detail and then describe economic applications. By now, there are many surveys of this literature, most notably the celebrated work by Roth and Sotomayor (1990) and

more recently by Abdulkadiroğlu and Sönmez (2013), Roth (2008a, b), Sönmez and Ünver (2009), Pathak (2015), Kojima and Troyan (2011), and Kojima (2015), and many others. Given the rich set of existing surveys, in this article, we try to balance between the basic models and recent applications. We also try to differentiate by choosing several specific topics which we regard as promising for further investigation.

The rest of this paper goes as follows. Section “Two-Sided Matching” presents the standard models of two-sided matching. In section “One-sided Matching,” we describe the models of one-sided matching. Section “Applications” discusses various applications. Section “Conclusion” concludes by discussing several future research directions.

Two-Sided Matching

In two-sided matching, there are two groups of participants. Each participant may be matched to a participant on the other side of the market (or remain unmatched), and she has preferences over these options. Two canonical examples are college admissions and firm-worker assignment. In the first example, students try to get into their “ideal” colleges, and colleges seek to admit their most preferred students. In the second example, workers look for their “dream” jobs, and firms attempt to fill the openings with their desired talents. In the presence of mutual interests and potential conflicts, of course, it is not generally possible to completely fulfill all participants’ desires. In the face of this constraint, it would be desirable if a procedure can help match the students (workers) with colleges (firms) in a fair and efficient manner. That is precisely what we will take up in the current section.

On the other hand, despite their close resemblance, there is at least one notable difference between the college admissions problem and the firm-worker assignment: the terms governing the match of a student and a college are almost always identical, so the preference of a student/college

depends only on the identity of her/its partners (if we ignore the differences in fellowship and other flexible terms). By comparison, labor contracts may vary much, with wage differences as a prominent example.¹ As such, a worker/firm not only cares about who he/it is matched with but also the contracting details as well. Because of this additional complication, the firm-worker assignment is naturally more involved than the college admissions problem.

We will begin by presenting the basic Gale and Shapley (1962) two-sided matching model in section “[Basic Two-Sided Matching Model](#),” with the college admissions problem as the leading example. In section “[Matching with Contracts](#),” we discuss the matching with contract model (Hatfield and Milgrom 2005), with the firm-worker assignment in mind.

Basic Two-Sided Matching Model

Adopting the language of Gale and Shapley (1962) and Roth (1985), we describe the basic model in terms of colleges and students. However, we note that it can be applied to any matching model where both sides of the markets have preferences, and the contracting details are standardized (e.g., medical residency matching).

There is a finite set I of students and a finite set C of colleges. A student can be matched to at most one college, while each college c has capacity q_c , i.e., the college can be matched to at most q_c students. Each student i has a strict preference $\succ_i$ over $C \cup \{\emptyset\}$, where $\emptyset$ denotes the outcome in which the student is unmatched, and each college has a strict preference $\succ_c$ over sets of students 2^I . For student i , we write $c_1 \succ_i c_2$ if and only if $c_1 \succ_i c_2$ or $c_1 = c_2$. Similarly, for college c , we write $J_1 \succ_c J_2$ if and only if $J_1 \succ_c J_2$ or $J_1 = J_2$. Note we implicitly assume a student/college only cares about his/its own match and that indifference does not occur.

Throughout the current section, we assume the preference of each college c is **responsive**, or the

relative desirability of sets of students does not depend on the composition of the current assignment of college c . More formally, $\succ_c$ is **responsive** if:

1. For any $J \subset I$ with $|J| < q_c$ and any $i \in I \setminus J$, $(J \cup i) \succ_c J \Leftrightarrow i \succ_c \emptyset$
2. For any $J \subset I$ with $|J| < q_c$ and any $i, j \in I \setminus J$, $(J \cup i) \succ_c (J \cup j) \Leftrightarrow i \succ_c j$. (To simplify notation, we denote a singleton set $\{s\}$ as s whenever there is no confusion.)

We write the set of all strict responsive preference profiles as:

$$\mathcal{R} = \{(\succ_i)_{i \in I \cup G} \mid \succ_c \text{ is responsive, } \forall c \in C\}.$$

A **matching** is a function $\mu : I \rightarrow C \cup \{\emptyset\}$. For each $c \in C$, we define $\mu(c) = \{i \in I \mid \mu(i) = c\}$. We say that a matching μ is **feasible** if $|\mu(c)| \leq q_c$ for all $c \in C$. Simply put, in a feasible matching, each college is matched with a set of students not exceeding its capacity. For the rest of the discussion, we only look at feasible matchings and will simply refer to them as matchings. Let $\mathcal{M}$ be the set of all (feasible) matchings.

As mentioned at the beginning of the section, our goal is to find a systematic procedure that can help match students with colleges in a fair and efficient manner. Before proceeding, we first make precise what we mean by “a systematic procedure” and “a fair and efficient manner.” Formally, a **(direct) mechanism** is a function that produces a (random) matching (outcome) for each preference profile, or $\varphi : \mathcal{R} \rightarrow \Delta \mathcal{M}$. For fairness and efficiency, one possible criterion is **stability**. Formally, a matching μ is **individually rational** if $\mu(i) \preceq_i \emptyset$, $\forall i \in I$ and $i \preceq_c \emptyset$, $\forall i \in \mu(c)$, $c \in C$. A matching μ is **blocked by a pair** $(i, c) \in I \times C$ if:

1. $c \succ_i \mu(i)$.
2. $|\mu(c)| < q_c$ and $i \succ_c \emptyset$ or $|\mu(c)| = q_c$ and $i \succ_c$ for some $j \in \mu(c)$.

We say a matching is **stable** if it is individually rational and not blocked by any pair; a mechanism is **stable** if it produces a stable matching for any preference profile in any realization. Intuitively, a

¹Note, however, terms may not vary substantially in some labor markets, especially in standardized entry-level markets. In such a case, the matching model in section “[Basic Two-Sided Matching Model](#)” may be appropriate.

matching is individually rational if the assignment of each student and college is acceptable (at least as good as being unmatched). A matching is not blocked by any pair if whenever a student prefers a college to his current assignment, either the student is not acceptable to the college or the college has reached its capacity, and it prefers every current student to the new student (note that we have implicitly made use of the assumption that preferences are responsive). A matching is stable if both conditions are satisfied. Stability is a desirable criterion for at least two reasons: to begin with, in an unstable mechanism, a student/college or a pair may want to deviate from the proposed outcome, which is problematic if we require voluntary participation. In addition, stability implies standard concepts of efficiency and fairness.

To see this, recall a matching μ is **(strongly) Pareto efficient** if there is no other matching ν such that $\nu(l) \succsim_{\mu}(l)$ for all $l \in I \cup C$ and $\nu(l) \succ_{\mu}(l)$ for some $l \in I \cup C$. Given a matching μ , a student i has **justified envy** toward student $j \in \mu(c)$, if $c \succ_{\mu}(i)$ and $i \succ_j j$. We say μ is **justified envy-free** if it is individually rational and no student has justified envy. Roughly speaking, (strong) Pareto efficiency says the assignment of a student or a college can only be improved at the expense of another student or college. Justified envy-freeness says if a student prefers the college of another student to his current matching, then it must be the case the college ranks the other students higher. If we take (strong) Pareto efficiency and envy-freeness as the natural efficiency and fairness requirement, then the following proposition shows they are both implied by stability.

Proposition 1 *If a matching μ is stable, then it is both (strongly) Pareto efficient and envy-free.*

Having established a desirable property of a matching (and thus a mechanism), one may wonder whether a stable matching can be achieved with every preference profile. Gale and Shapley (1962) gave a positive answer with the construction of the following mechanism. Since then, the mechanism and its variations have played important roles in real-life matching markets.

(Student-Proposing) Deferred Acceptance Algorithm:

- Step 1: Each student applies to his most preferred college. Each college tentatively keeps its most preferred acceptable students up to the capacity (on hold) and rejects all others.

In general, for any $t = 1, 2, \dots$

- Step t : Each student who is not currently on hold applies to his next preferred acceptable college (he does not make an application if there is none). Each college considers all new applicants, together with the students on hold, and tentatively keeps its most preferred acceptable students up to the capacity while rejecting all others.

The algorithm terminates when there is no new application. (Clearly, it terminates in a finite number of steps because the number of students and colleges are both finite.)

Theorem 1 (Theorem 1 in Gale and Shapley 1962) *For any (strict, responsive) preference profile, the (student-proposing) deferred acceptance algorithm gives a stable matching. In other words, it is a stable mechanism.*

Given the student-proposing deferred acceptance algorithm, one may naturally imagine a corresponding version of the college-proposing deferred acceptance algorithm. Such a mechanism does exist and is also stable. In fact, the set of stable matchings is not necessarily a singleton set, and the student-proposing and college-proposing deferred acceptance algorithms can give rise to different stable matchings. Given the (potential) multiplicity of stable matchings, one may wonder which one to implement in practice. The following proposition suggests the answer depends on our evaluation of the relative welfare of the two sides of the markets.

Proposition 2 (Theorem 2* in Roth 1985)

1. *There exists a student-optimal stable matching, i.e., a stable matching that every student likes at least as well as any other stable matching. Moreover, the student-proposing deferred acceptance algorithm always yields the student-optimal stable matching.*

2. *There exists a college-optimal stable matching, i.e., a stable matching that every college likes at least as well as any other stable matching. Moreover, the college-proposing deferred acceptance algorithm always yields the college-optimal stable matching.*
3. *The student-optimal stable matching is the least preferred stable matching for each college. Likewise, the college-optimal stable matching is the least preferred stable matching for each student.*

Roughly speaking, Proposition 2 says that despite the competition among students/colleges, there is considerable coincidence of interests on either side of the market if we restrict to stable matchings. By comparison, the interests of the two sides of the markets are not always aligned. In fact, they are almost opposite if we restrict our attention to stable matchings.

Another important concern for mechanisms is incentives. Apart from ease of implementation, an incentive compatible mechanism induces students and colleges to reveal their preferences truthfully. Such a property is necessary for our fairness and efficiency criteria to be well-grounded.² We say a mechanism is **strategy proof** if it is a weakly dominant strategy for every player to (always) report his/her true preferences. With this definition, we have the following result for incentive properties of the student-proposing deferred acceptance algorithm (and any stable mechanism).

Theorem 2 (Theorem 5* in Roth 1985) *The student-proposing deferred acceptance algorithm is strategy proof for students. However (when colleges have responsive preferences), no stable mechanism is strategy proof for colleges.*

Simply put, Theorem 2 suggests that the student-proposing deferred acceptance algorithm

is “safe” to play for students: it is always optimal for students to simply report their true preferences. The rough intuition is that because of “deferred” acceptance, a student stands nothing to lose by applying to his most preferred (remaining) college in each step. Furthermore, even though it is not always in the colleges’ interests to report truthfully, the problem is not confined to the particular mechanism, but stable mechanisms in general. In other words, incentive compatibility (on the college side) is not compatible with stability.

To see how colleges may gain by misreporting under the student-proposing deferred acceptance algorithm, consider the following example:

Example 1 There are two students i_1 and i_2 and two colleges c_1 and c_2 . Each college has capacity 1 ($q_{c_1} = q_{c_2} = 1$), and the preferences are as follows: (When denoting an agent’s preference, we list her acceptable choices in order of her preference. Similar notation is used throughout the paper.)

$$\begin{array}{ll} \succ_{i_1} : c_1, c_2 & \succ_{i_2} : c_2, c_1 \\ \succ_{c_1} : i_2, i_1 & \succ_{c_2} : i_1, i_2 \end{array}$$

It is easy to see the outcome of the student-proposing deferred acceptance algorithm; under the true preferences is:

$$\begin{pmatrix} i_1 & i_2 \\ c_1 & c_2 \end{pmatrix}$$

Now suppose college 1 misreports by stating that only student 2 is acceptable, while all other agents continue to report their true preferences. Then, the outcome of the student-proposing deferred acceptance algorithm under this preference profile is:

$$\begin{pmatrix} i_1 & i_2 \\ c_2 & c_1 \end{pmatrix}$$

Thus, college 1 strictly benefits from the misreport.

Matching with Contracts

This section introduces the model of “[Matching with Contracts](#)” due to Hatfield and Milgrom

²Chen and Sonmez (2006) show evidence in lab experiments that strategy-proof school choice mechanisms indeed induce true preferences more often than those without truthful revelation. Nevertheless, there is some evidence in real-life matching markets that some agents still misreport even in strategy-proof mechanisms. See, for instance Rees-Jones (2017) and Hassidim et al. (2017).

(2005). It is a generalization of the basic model described in section “[Basic Two-Sided Matching Model](#).” It also incorporates, as special cases, some other matching/auction models in the existing literature. The primary example we have in mind is labor market matching, and so we describe the model in terms of workers and firms. However, it can be applied in any setting where contract terms play an important role.

There is a finite set I of workers, a finite set F of firms, and a finite set of contracts X . Each contract $x \in X$ is bilateral, so it is associated with one worker $x_I \in I$ and one firm $x_F \in F$. For instance, in the labor market matching model, a contract specifies a firm, a worker, and a wage. So we have $X = I \times F \times W$, where W is a (finite) set of possible wages.

Each worker i can sign at most one contract, and his preferences over possible contracts (plus the outcome in which she signs no contract, i.e., the empty set $\emptyset$, which we sometimes refer to as the null contract) are described by the strict total order $\succ_i$. We say a contract x is **acceptable** for worker i if $x \succ_i \emptyset$. With workers’ preferences well defined, we know a worker’s choice when faced with a set of contracts. Formally, given a set of contracts $X' \subset X$, define worker i ’s chosen set $C_i(X')$ of contracts as follows:

$$C_i(X') = \begin{cases} \emptyset & \text{if } \{x \in X' \mid x_I = i, x \succ_i \emptyset\} = \emptyset \\ \max_{\succ_i} \{x \in X' \mid x_I = i\} & \text{otherwise} \end{cases}$$

In other words, a worker’s chosen set is simply the most preferred, *acceptable* contract from those that are available. If none is acceptable, then his choice is the null contract (note that a worker’s chosen set is either a singleton or an empty set).

On the other hand, each firm f can sign multiple contracts, so its preferences are over sets of contracts. Let firm f ’s preference be described by $\succ_f$, a strict total order over 2^X . Given a set of contracts $X' \subset X$, we can similarly define firm f ’s chosen set $C_f(X')$. Notice that a firm can sign at most one contract with any given worker, so for all $f \in F$, $X' \subset X$, and $x, x' \in C_f(X')$, if $x \neq x'$, then $x_I \neq x'_I$.

Given $X' \subset X$ and the chosen set of each worker, we can define the contracts chosen by

the worker side as $C_I(X') = \cup_{i \in I} C_i(X')$. The remaining offers in X' are the rejected set by the worker side: $R_I(X') = X' - C_I(X')$. Similarly, the chosen and rejected sets by the firm side are denoted by $C_F(X') = \cup_{f \in F} C_f(X')$ and $R_F(X') = X' - C_F(X')$.

If we imagine some sort of stable matchings similar to the one in the basic model, the chosen and rejected sets will prove useful. Intuitively, given a starting set of contracts, the rejected sets (by either side) cannot be in any stable allocation, so we know, at the very least, a stable allocation is a fixed point of the “chosen set (by either side)” operator. Moreover, we will need to require that there is no coalition of workers and firms who prefer to sign contracts among themselves rather than follow the prescribed allocation.

To make this precise, we present the following definition of a stable allocation: a set of contracts $X' \in X$ is a **stable allocation** if:

1. $C_I(X') = C_F(X') = X'$.
2. There exists no firm f and a set of contracts $X'' \neq C_f(X')$ such that $X'' = C_f(X' \cup X'') \subset C_I(X' \cup X'')$.

Intuitively, the first requirement corresponds to “individual rationality.” The second requirement says there cannot be an alternative set of contracts such that a particular firm and its matched workers all prefer, which is similar to the “no blocking” condition we discussed in the basic model.

In order to guarantee the existence of a stable allocation, we need an additional restriction on the preferences of firms: substitutability. Elements of a set of contracts X are **substitutes** for firm f ; if $\forall X' \subset X'' \subset X$, we have $R_f(X') \subset R_f(X'')$. In words, the restriction says if a particular firm is faced with a larger choice set, it has to reject (weakly) more contracts.

Now we are ready to present the existence result, based on an iterative algorithm. As it turns out, the iteration we shall apply is on the product set $X \times X$ instead of the set of all contracts X . Similar to most iteration procedures, we hope to obtain monotonicity. For monotonicity to be well defined, a partial order on $X \times X$ is needed. We define it as follows:

given $X_I, X'_I, X_F, X'_F \subset X$, $((X_I, X_F) \geq (X'_I, X'_F)) \Leftrightarrow (X'_I \subset X_I \text{ and } X_F \subset X'_F)$.

With the order $\geq$, and given a starting set (X_I, X_F) , we can define the **generalized deferred acceptance algorithm** as the iterated applications of the function $F : X \times X \rightarrow X \times X$, defined by:

$$\begin{aligned} F_1(X') &= X - R_F(X') \\ F_2(X') &= X - R_I(X') \\ F(X_I, X_F) &= (F_1(X_F), F_2(F_1(X_I))). \end{aligned}$$

With the generalized deferred acceptance algorithm in hand, the following theorem and proposition tell us how they can direct us to find stable allocations and shed light on the trade-off in welfare between the workers and firms.

Theorem 3 (Theorem 3 in Hatfield and Milgrom 2005) *Suppose elements of X are substitutes for all the firms. Then,*

1. *The set of fixed points of F on $X \times X$ includes a smallest element $(\underline{X}_I, \underline{X}_F)$ and a largest element $(\bar{X}_I, \bar{X}_F)$.*
2. *Starting at $(X_I, X_F) = (X, \emptyset)$, the generalized deferred acceptance algorithm converges monotonically to the largest fixed point $(\bar{X}_I, \bar{X}_F) = \sup\{(X', X'') \mid F(X', X'') \geq (X', X'')\}$.*
3. *Starting at $(X_I, X_F) = (\emptyset, X)$, the generalized deferred acceptance algorithm converges monotonically to the smallest fixed point $(\underline{X}_I, \underline{X}_F) = \inf\{(X', X'') \mid F(X', X'') \leq (X', X'')\}$.*

Proposition 3 (Theorem 4 in Hatfield and Milgrom 2005) *Suppose elements of X are substitutes for all the firms, then $\bar{X}_I \cap \bar{X}_F$ and $\underline{X}_I \cap \underline{X}_F$ are both stable allocations. Moreover, every worker likes $\bar{X}_I \cap \bar{X}_F$ at least as well as any other stable allocation, and every firm likes any other stable allocation at least as well as $\bar{X}_I \cap \bar{X}_F$. Similarly, every firm likes $\underline{X}_I \cap \underline{X}_F$ at least as well as any other stable allocation and every worker likes any other stable allocation at least as well as $\underline{X}_I \cap \underline{X}_F$.*

To understand Theorem 3 and Proposition 3, consider first the generalized deferred acceptance algorithm starting with the contract tuple $(X, \emptyset)$. Here, workers first choose from the set of all available contracts, which is very similar to the

student-proposing deferred acceptance algorithm (where students first propose to their most preferred colleges). Accordingly, we land at the worker-optimal stable allocation. Similarly, the generalized deferred acceptance algorithm starting with the contract tuple $(\emptyset, X)$ gives us the firm-optimal stable allocation. For Proposition 3, the first part of why $\bar{X}_I \cap \bar{X}_F$ and $\underline{X}_I \cap \underline{X}_F$ is stable that needs some additional reasoning (which we shall not give here), but the second part incorporates the insight from the basic model: if we focus on stable allocations, there is little conflict of interests among agents on the same side of the market, but there is substantial conflict of interests between the worker side and the firm side.

Similar to the basic model, incentives are also an important concern in the matching with contracts model, but we shall omit it here due to space constraint.

One-Sided Matching

Similar to two-sided matching, in a one-sided matching problem, there are still two groups. Nevertheless, one side has no (intrinsic) preference for the other side and is simply “objects to be consumed.” Two typical examples are house allocation and kidney exchange. In each problem, each individual tries to obtain her most preferred object and, if she has an initial endowment, may obtain her desired object through exchange. Our goal is to find a systematic procedure to allocate the houses (kidneys) to tenants (patients) in a fair and efficient way.

Even if we ignore the contracting details such as side payments (as in two-sided matching), complications arise depending on the structure of initial endowments. The first case we consider is pure exchange: each individual brings an item to the market and tries to achieve an efficient outcome that everyone is willing to accept. The next case is pure allocation: nobody has anything a priori, and there are a number of items to be allocated to the individuals. A more involved case is allocation with existing owners: some individuals come to the market with an item, while others do not, and there are new items

available for allocation. In this situation, we want to find an allocation that appeals to both the existing owners and the newcomers. We will present the three different models in sections “[House Exchange](#),” “[House Allocation with No Existing Owner](#),” and “[House Allocation with Existing Owners](#).” To highlight the similarities and differences, we describe all three models in the language of housing markets.

House Exchange

As the name suggests, in a house exchange problem, each individual is initially endowed with a house, and there is no new house available. Thus, the only possible reallocation is through exchange between the agents. The model was first described by Shapley and Scarf (1974). (To better conform with the two-sided matching model, our notation will be different from theirs. However, the model and the main results follow theirs.)

There is a finite set H of houses and a finite set I of house owners. Initially, individual i is the owner of house h_i , with $|H| = |I|$. Therefore, everyone is endowed with exactly one house to begin with, and there is no leftover house. Each agent i demands exactly one house and has a strict preference $\succ_i$ over H (we assume all houses are acceptable, though it is enough to assume agent i 's initial endowment h_i is acceptable). Similar to two-sided matching, we write $h_1 \succ_i h_2$ if and only if $h_1 \succ_i h_2$ or $h_1 = h_2$. We also implicitly assume a house owner only cares about her own assignment, and indifferences between houses do not occur. We write the set of all strict preference profiles as $\mathcal{R} = \{(\succ_i)_{i \in I}\}$. A matching is a one-to-one correspondence $\mu : I \rightarrow H$ (or equivalently represented as a permutation on $\{1, 2, \dots, |I|\}$). Note that there is no feasibility concern once we restrict to one-to-one correspondences. Let $\mathcal{M}$ be the set of all matchings. A (direct) mechanism in the one-sided matching model is a function $\varphi : \mathcal{R} \rightarrow \Delta\mathcal{M}$.

Given one-sided preference and no priority, one possible criterion in the current setup is the **(strong) core** standard in cooperative games. Given $(\succ_i)_{i \in I}$, recall a matching μ is a **(strong) core** allocation if there is no other matching $\mu' \in \mathcal{M}$ and a subset $S \subset I$ such that:

1. $\mu'(i) \in \{h_j\}_{j \in S}, \forall i \in S$
2. $\mu'(i) \succ_i \mu(i), \forall i \in S$
3. $\mu'(j) \succ_j \mu(j)$ for some $j \in S$

In words, a matching is a strong core allocation if no subgroup can come up with an allocation using only their endowments that every group member weakly prefers to the original allocation, and some group member strictly prefers. Similar to stability, (strong) core reflects the idea of voluntary participation. Moreover, if we take S to be the singleton set i and the whole set I , we see the definition implies individual rationality and strong Pareto efficiency (here the outside option of being unmatched is modified to be the initial endowment in the definition of individual rationality).

Proposition 4 *If a matching μ is a (strong) core allocation, then it is both individually rational and (strongly) Pareto efficient.*

Notice that with one-sided preference and no priority list, fairness is not too much of a concern once we respect individual rationality. Given the desirable criterion, our goal is to find a strong core allocation (if any) for any preference profile. Fortunately, David Gale (described in Shapley and Scarf (1974)) gave a satisfactory answer with the construction of the top trading cycle algorithm.

Top Trading Cycle Algorithm:

- Step 1: Each agent points to her most preferred house and each house points to its owner. Remove all agents and houses in a cycle (at least one cycle exists). For any agent removed, assign her the house she points to.

In general, for any $t = 1, 2, \dots$

- Step t : Each remaining agent points to her most preferred house left, and each house points to its owner (who must still be in the mechanism). Remove all agents and houses in a cycle (at least one cycle exists). For any agent removed, assign her the house she points to.

The algorithm terminates when there is no agent or house left. (An equal number (greater

than or equal to one) of agents and houses is removed in each step, so the mechanism must terminate in a finite number of steps.)

Theorem 4 (Theorem in Shapley and Scarf 1974 and Theorem 2 in Roth and Postlewaite 1977) *For any strict preference profile, the top trading cycle algorithm gives the unique strong core allocation.*

Intuitively, the top trading cycle gives a strong core allocation roughly because (given the cycles removed earlier), the group of agents removed in each step receive their best available houses.

As in two-sided markets, another important concern in house exchange is incentives. After all, it is only with truthful revelation that the strong core requirement has important welfare implications. Fortunately, the following theorem says the top trading cycle algorithm has good incentive properties.

Theorem 5 (Theorem in Roth 1982) *The top trading cycle mechanism is strategy proof.*

To obtain intuition for Theorem 5, recall truthful revelation is a weakly dominant strategy for students in the student-proposing deferred acceptance algorithm. The insight here is somewhat similar: a house will not leave the market unless its (initial) owner gets her most preferred (remaining) house. Therefore, an agent will not “lose” a house unless she cannot get it anyway. Hence, she may as well report her most preferred (remaining) house in each step.

Given the relative complexity of the top trading cycle algorithm, we end the section with an example:

Example 2 There are four house owners i_1, i_2, i_3, i_4 , with their respective (initial) houses h_1, h_2, h_3, h_4 . The preferences of the house owners are as follows:

$$\begin{array}{ll} \succ_{i_1} : h_2, h_3, h_4, h_1 & \succ_{i_2} : h_2, h_1, h_3, h_4 \\ \succ_{i_3} : h_2, h_1, h_3, h_4 & \succ_{i_4} : h_2, h_1, h_3, h_4 \end{array}$$

Step 1: All agents point to h_2 and the houses point to their respective owners. There is a one cycle: i_2 is assigned h_2 .

Step 2: i_1 points to h_3 , i_3 and i_4 point to h_1 , and the three remaining houses point to their respective owners. There are two cycles: i_1 is assigned h_3 and i_3 is assigned h_1 .

Step 3: i_4 points to h_4 , which points backward. There is a one cycle: i_4 is assigned h_4 .

It follows that the outcome of the top trading cycle algorithm is:

$$\begin{pmatrix} i_1 & i_2 & i_3 & i_4 \\ h_3 & h_2 & h_1 & h_4 \end{pmatrix}$$

House Allocation with No Existing Owner

In a pure house allocation problem, no agent has a house to begin with, and there are a number of houses to be allocated. A similar problem was first studied by Hylland and Zeckhauser (1979). The model we present here is a simplified version of theirs.

There is a finite set H of empty houses to be allocated among a finite set I of agents. Each agent i demands exactly one house and has a strict preference $\succ_i$ over H . (We assume all houses are acceptable.) It may be the case that $|H| > |I|$, $|H| < |I|$, or $|H| = |I|$, so the houses may be in oversupply and undersupply or just balance with the number of agents. Similar to the house exchange problem, we write $h_1 \succsim_i h_2$ if and only if $h_1 \succ_i h_2$ or $h_1 = h_2$. Implicit in the assumption is that an agent only cares about her own assignment and that indifferences between houses do not occur. We write the set of all strict preference profiles as $\mathcal{R} = \{(\succ_i)_{i \in I}\}$. A matching is a one-to-one function $\mu : I \rightarrow H \cup \emptyset$ (as the number of houses and agents need not balance, in general, a matching here is not bijective and cannot be equivalently represented as a permutation on $\{1, 2, \dots, |I|\}$). Let $\mathcal{M}$ be the set of all matchings. A (direct) mechanism is a function $\varphi : \mathcal{R} \rightarrow \Delta\mathcal{M}$.

Given the lack of existing owners, the primary criterion here is (strong) Pareto efficiency. Nevertheless, randomization may be particularly useful in the current setup if there are fairness concerns. Once we introduce randomization, at least two versions of (strong) Pareto efficiency arise.

A mechanism is **ex ante Pareto efficient** if its assignment of lotteries is Pareto efficient relative to agents' preferences over lotteries. By comparison, a mechanism is **ex post Pareto efficient** if its final allocation is Pareto efficient given any strict preference profile. It can be readily shown that ex ante Pareto efficiency implies ex post Pareto efficiency but not vice versa.³

Before introducing a desirable algorithm, one more definition is needed: a **(rank) ordering** is a permutation of I , or a one-to-one correspondence $\sigma : \{1, 2, \dots, |I|\} \rightarrow I$. The following mechanism and its variations are widely used in real-life house allocation problems:

Serial Dictatorship Algorithm:

- Step 0: Fix a rank ordering σ .
- Step 1: Assign $\sigma(1)$ her most preferred house.

In general, for any $t = 1, 2, \dots$

- Step t : Assign $\sigma(t)$ her most preferred remaining house.

The algorithm terminates when there is no agent or house left. If there are still agents left, then they are not assigned a house. (Given each step reduces the number of agents and houses both by 1, the algorithm must terminate in a finite number of steps.)

Intuitively, the mechanism works as if an agent is the dictator when it is her turn to choose. At that time, she picks her most preferred house out of those available (note that the agent does not care about the allocation of any other agent). As mentioned earlier, randomization is often introduced when implementing the mechanism in practice. This can be done by modifying Step 0 as follows:

- Step 0: Pick a rank ordering uniformly at random from the set of all rank orderings.

The resulting mechanism is called **random serial dictatorship**. The following theorem gives the desirable properties of random serial dictatorship.

Theorem 6 (Variation of Lemma 1 in Abdulkadiroğlu and Sönmez 1998) *The random serial dictatorship algorithm is ex post Pareto efficient. Moreover, two agents with the same preferences receive the same random allocation to each other.*

In other words, the random serial dictatorship mechanism has decent fairness and efficiency properties. Intuitively, ex post efficiency is achieved because an agent is made as well off as possible given the allocation of the agents in earlier steps. (This is true in every realization, so “ex post” with randomization.) The second part of this theorem describes a fairness property of this mechanism, and it follows immediately from uniform randomization over rank orderings used in Step 0’ of the algorithm.

The following theorem says that random serial dictatorship also has good incentive properties.

Theorem 7 The random serial dictatorship mechanism is strategy proof.

The main insight of Theorem 7 is that each agent is essentially the dictator when it comes to her turn, so she cannot gain by misreporting.

Unfortunately, even random serial dictatorship is not without its own problems. For one, the mechanism is not ex ante Pareto efficient. Some research has been done to tackle the problem. Nevertheless, it is found that ex ante Pareto efficiency and fairness (as defined in Theorem 6) are incompatible with strategy proofness. (See Bogomolnaia and Moulin (2001).) Partly because of this, random serial dictatorship is still probably among the most popular mechanisms when it comes to real-life object allocation.

House Allocation with Existing Owners

Given the discussions of sections “House Exchange” and “House Allocation with No Existing Owner,” one may imagine a situation where existing house owners and new entrants coexist. A few more specific real-life examples are college dorm allocations and office

³Ex-ante Pareto efficiency implies ex-post Pareto efficiency because if any final allocation resulting from a lottery is not ex-post Pareto efficient, then the lottery can be improved by replacing the particular allocation with a more efficient one, implying that the lottery is not ex-ante Pareto efficient.

assignment. The problem was first studied by Abdulkadiroğlu and Sönmez (1999).

There are a finite set of houses H and a finite set of agents I . Of all the houses in H , a subset H_O is currently occupied, each belonging to a distinct member of the existing house owners $I_E \subset I$ (so $|H_O| = |I_E|$). The remaining houses $H_V = H - H_O$ are currently vacant and can be freely allocated. The remaining agents $I_N = I - I_E$ are new entrants and do not have a house. Each agent $i \in I$ demands exactly one house and has a strict preference $\succ_i$ over H . (We assume for simplicity that all houses are acceptable for all the agents.) We write $h_1 \succ_i h_2$ if and only if $h_1 \succ_i h_2$ or $h_1 = h_2$. Implicit in the assumption is that an agent only cares about her own assignment and that indifferences between houses do not occur. We write the set of all strict preference profiles as $\mathcal{R} = \{(\succ_i)_{i \in I}\}$. A matching is a one-to-one function $\mu : I \rightarrow H \cup \emptyset$ (Similar to the pure house allocation problem, a matching here need not be bijective.) Let $\mathcal{M}$ be the set of all matchings. A (direct) mechanism is a function $\varphi : \mathcal{R} \rightarrow \Delta\mathcal{M}$.

Given the presence of both existing house owners and new entrants, one possible desirable criterion of a matching is Pareto efficiency. In light of the top trading cycle and serial dictatorship algorithms of previous sections, we have the following two generalizations as natural candidates. Indeed, Abdulkadiroğlu and Sönmez (1999) show for any preference profile; outcomes from these two mechanisms coincide and satisfy Pareto efficiency.

(Generalized) Top Trading Cycle Algorithm:

- Step 0: Fix a rank ordering σ .
- Step 1: Define the set of available houses to be the vacant houses (H_V). Each agent points to her most preferred house. Each occupied house points to its owner, and each available house points to $\sigma(1)$. Remove all agents and houses in a cycle (at least one cycle exists). For any agent removed, assign her the house she points to.

In general, for any $t = 1, 2, \dots$

- Step t : Update the set of available houses to be the current vacant houses. Each remaining

agent points to her most preferred house left. Each remaining occupied house points to its owner, and each available house points to the remaining agent with the highest priority ($\sigma(j)$, where j is the smallest among the remaining agents). Remove all agents and houses in a cycle (at least one cycle exists). For any agent removed, assign her the house she points to.

The algorithm terminates when there is no agent or house left. If there are still agents left, then they are not assigned a house. Since each step reduces the number of agents and houses both by at least 1, the algorithm must terminate in a finite number of steps.

You Request My House-I Get Your Turn (YRMH-IGYT) Algorithm:

- Step 0: Fix a rank ordering σ .
- Step 1: Agent $\sigma(1)$ points to her most preferred house. If the house she points to is vacant (in H_V) or her own house, she is assigned the house she points to. Otherwise, modify σ so that the owner is at the top of the list (the other relative orderings unchanged) and proceed to the next step.

In general, for any $t = 1, 2, \dots$

- Step t : The remaining agent with the highest priority ($\sigma(j)$, where j is the smallest among the remaining agents) points to her most preferred house. If the house she points to is *currently* vacant (which may or may not be in H_V) or her own house, she is assigned the house she points to. If the house she points to is occupied by another remaining agent, modify σ so that the owner is at the top of the list (the other relative orderings unchanged). At this point, if a loop forms (no house is assigned in the process where the rank ordering is back to an earlier one), every agent is assigned the house she points to. Otherwise, proceed to the next step.

The algorithm terminates when there is no agent or house left. If there are still agents left, then they are not assigned a house. (It can be shown the algorithm terminates in a finite number of steps.)

Intuitively, the (generalized) top trading cycle algorithm is a direct generalization of top trading cycles in section “[House Exchange](#),” with all remaining vacant houses pointing to the remaining agent with the highest priority. On the other hand, YRMH-IGYT is a direct generalization of serial dictatorship in section “[House Allocation with No Existing Owner](#),” with the added twist that the owner is granted the opportunity to choose before her house is gone.

Theorem 8 (Theorem 3 in Abdulkadiroğlu and Sönmez 1999) *Given a rank ordering σ and for any (strict) preference profile, the YRMH-IGYT algorithm yields the same matching as the generalized top trading cycle algorithm.*

Theorem 9 (Propositions 1 and 2 in Abdulkadiroğlu and Sönmez 1999) *Given a rank ordering σ and for any (strict) preference profile, the matching given by YRMH-IGYT and generalized top trading cycle algorithms is strongly Pareto efficient.*

Moreover, the following theorem reveals that incentives do not pose a problem either.

Theorem 10 (Theorem 1 in Abdulkadiroğlu and Sönmez 1999) *For any rank ordering σ , both the YRMH-IGYT and the generalized top trading cycle algorithms are strategy proof.*

Given the close relationships, the intuitions of Theorems 9 and 10 are very similar to the counterparts of top trading cycle and serial dictatorship algorithms (Theorems 4 through 7).

To illustrate the two algorithms and the main insights of Theorem 8, we conclude the section with an example.

Example 3 There are four agents and three houses. Agents i_1 and i_2 are current house owners, with their respective houses h_1 and h_2 . Agents i_3 and i_4 are new entrants. House h_3 is currently available. There preferences of the agents are as follows:

$$\succ_{i_1} : h_3, h_2, h_1 \quad \succ_{i_2} : \succ_{i_3} : \succ_{i_4} : h_1, h_3, h_2$$

Fix the rank ordering $\sigma = (3, 4, 1, 2)$.

Generalized Top Trading Cycles:

Step 1: i_1 points h_3 , and the remaining agents point to h_1 . h_1 points to i_1 , h_2 points to i_2 , and h_3 points to i_3 . There are two cycles: i_1 is assigned h_3 and i_3 is assigned h_1 .

Step 2: i_2 and i_4 both point to h_2 and h_2 points to i_2 . There is a one cycle: i_2 is assigned h_2 .

It follows that the outcome of the generalized top trading cycle algorithm is:

$$\begin{pmatrix} i_1 & i_2 & i_3 & i_4 \\ h_3 & h_2 & h_1 & \emptyset \end{pmatrix}$$

YRMH-IGYT:

Step 1: i_3 points to h_1 , which is currently occupied. The ranking ordering σ is modified to $(1, 3, 4, 2)$.

Step 2: i_1 points to h_3 , which is currently vacant and i_1 is assigned h_3 .

Step 3: i_3 points to h_1 , which is currently vacant and i_3 is assigned h_1 .

Step 4: i_4 points to h_2 , which is currently occupied. The ranking ordering σ is modified to $(1, 3, 2, 4)$.

Step 5: i_2 points to h_2 , which is her own house and i_2 is assigned h_2 .

It follows that the outcome of the YRMH-IGYT algorithm is also:

$$\begin{pmatrix} i_1 & i_2 & i_3 & i_4 \\ h_3 & h_2 & h_1 & \emptyset \end{pmatrix}$$

Applications

The theories described in the previous sections have found applications in a wide variety of areas. While we are not able to survey all of them extensively, we have selected some of the most prominent examples in order to highlight how the theory can be utilized. We begin by discussing the following topics:

1. Medical Residency Matching (most closely related to the model in section “[Basic Two-Sided Matching Model](#)”)

2. Kidney Exchange (most closely related to the model in section “[House Allocation with Existing Owners](#)”)
3. School Choice (most closely related to the models in sections “[Basic Two-Sided Matching Model](#)” and “[House Allocation with Existing Owners](#)”)

Then, building on an understanding of the problems encountered when applying the tools to the situations above, we transition to a relatively new area of matching theory called “matching with constraints.”

Medical Residency Matching

Among the most common application of two-sided matching algorithms is the medical residency programs. In 2016, roughly 43,000 medical school graduates registered for the National Resident Match Program (NRMP), where students are matched to teaching hospitals through a variant of a deferred acceptance algorithm (for more detailed statistics, one can visit <http://www.nrmp.org/match-data/main-residency-match-data/>). This service has been in operation since 1952, and its longevity is ascribed to the fact that the matchings produced are stable (Roth 1984; Roth and Sotomayer 1990).

What makes NRMP’s matching problem complex, though, is the existence of “couples.” While some students apply independently and rank their preferences accordingly, individuals who have a significant other in the residency match program are allowed to apply together as couples so that they can work in areas close to one another. In this context, stability requires there be no coalition of students and hospitals who prefer to match among themselves than follow the prescribed matching.⁴ The presence of couples who submit joint preference lists complicates the problem significantly as

stability is not guaranteed. Take the following example (from Roth 1984):

Example 4 There are four medical school graduates i_1, i_2, i_3 , and i_4 and four hospitals h_1, h_2, h_3 , and h_4 each with capacity 1 ($q_{h_1} = q_{h_2} = q_{h_3} = q_{h_4} = 1$). (i_1, i_2) and (i_3, i_4) are couples with preferences over *ordered pairs* of hospitals. The exact preferences of the couples and the hospitals are as follows:

$$\begin{aligned}
 \succ_{(i_1, i_2)} &: (h_1, h_2), (h_4, h_1), (h_4, h_3), (h_4, h_2), (h_1, h_4), \\
 &(h_1, h_3), (h_3, h_4), (h_3, h_1), (h_3, h_2), (h_2, h_3), (h_2, h_4), \\
 &(h_2, h_1) \\
 \succ_{(i_3, i_4)} &: (h_4, h_2), (h_4, h_3), (h_4, h_1), (h_3, h_1), \\
 &(h_3, h_1), (h_3, h_2), (h_3, h_4), (h_2, h_1), (h_2, h_3), (h_1, h_2), \\
 &(h_1, h_4), (h_1, h_3) \\
 \succ_{h_1} &: i_4, i_2, i_1, i_3 \succ_{h_2} : i_4, i_3, i_2, i_1 \\
 \succ_{h_3} &: i_2, i_3, i_1, i_4 \succ_{h_4} : i_2, i_4, i_1, i_3
 \end{aligned}$$

It is straightforward, if tedious, to check that no stable matching exists in this example. Given that an increasing number of medical students marry other medical students, it would seem then that finding a stable matching for NRMP would be impossible. Even determining, in a given instance, whether a stable matching exists is a computationally hard problem.⁵ As a result, to replace the original method, Roth and Peranson (1999) proposed a heuristic modification of the deferred acceptance algorithm in place to accommodate couples’ preferences. Although this algorithm is not guaranteed to always produce a match that is stable with respect to the reported preferences, it has done so in almost all instances.

Why does the algorithm in NRMP find a stable matching despite the theoretical possibility of nonexistence? Kojima et al. (2013) show that in a setting where applicant preferences are drawn independently from a distribution, as the size of the market increases and the proportion of couples approaches 0, the Roth and Peranson algorithm terminates in a stable matching with high probability. Thus, one of the reasons the NRMP algo-

⁴The main difference of this definition from the one in the basic model of section “[Basic Two-Sided Matching Model](#)” is that we consider a coalition composed of a couple of doctors and two hospitals each of which seeks to match with a member of the couple. See Roth (1984) for detail.

⁵More precisely, this problem is in the class of “NP-hard” problems. NP-hardness is a notion in computational complexity theory describing the complexity of computation, which we will not describe in detail here.

algorithm finds stable matchings in most cases may be because the size of NRMP is large, while the proportion of couples in the market is small, roughly between 5% and 10%. By contrast, Biro and Klijn (2013) and Ashlagi et al. (2014) have shown, in separate settings, that as the proportion of couples increases, this algorithm frequently fails to terminate in a stable matching. This may be important given that residency matching is not the only environment with a “couples” issue. In other such settings, couples could make up much more of the market. (Biro and Klijn (2013) provide the example of assigning high school teachers in Hungary to majors, where almost all teachers need to be assigned to two majors; in this setting, the percentage of “couples” is nearly 100%).

Given these difficulties in the “couples” problem, Nguyen and Vohra (2017) propose an alternative approach. They allow for perturbations of hospital capacities to find a “nearby” instance of the matching problem that is guaranteed to have a stable matching. They find that the necessary perturbations are small, especially when hospital/school/firm capacities are large. Specifically, given capacities q_h for each hospital h , there is a redistribution of the slots, q'_h , satisfying $|q_h - q'_h| \leq 2$ for all hospitals h and $\sum q_h \leq \sum q'_h \leq \sum q_h + 4$. Thus, the perturbations change the capacity of each individual hospital by at most 2 and increase the total number of positions in hospitals by no more than 4 while never decreasing it.⁶

The complication surrounding matching with couples turns out to be a specific instance of a more general issue economists have sought to understand: matching with complementarities. In two-sided matching markets, substitutability of agent preferences (see section “[Matching with Contracts](#)”), i.e., the lack of complementarity, is

“necessary” for guaranteeing the existence of a stable matching (see Hatfield and Kojima (2008) and Sönmez and Ünver (2010) for formal statements). The existence of couples leads to a violation of substitutability because a pair of positions close to each other works as complements for the couple. Recent research by Che et al. (2017) and Azevedo and Hatfield (2017) examine matching with complementarities in large market settings with a continuum of agents. They have found positive results describing sufficient conditions for the existence of stable matchings.

Kidney Exchange

The application of matching theory to kidney exchange has been discussed often and is quite thorough, so we will be relatively brief in our exposition. For an extensive survey, we refer the reader to Sönmez and Ünver (2011). In the kidney “market” (using the term loosely), the National Organ Transplant Act of 1984 made it illegal to buy or sell a kidney in the USA. Similar legal prohibitions are nearly universal around the globe. Thus, donation is the only viable option for kidney transplantation for most patients.

The initial foundational contribution to kidney exchange came with Roth et al. (2004). They used a variation of the Shapley-Scarf house exchange model (section “[House Exchange](#)”) to represent the kidney exchange market. In their model, agents enter in pairs composed of a patient and his potential donor. Applying the top trading cycle (TTC) mechanism where potential donors substitute “houses” of the original Shapley-Scarf model, one can produce a matching between donors and patients in a Pareto-efficient and strategy-proof way.

The way economists model kidney exchange has progressed as we now know many ways in which the assumptions in the original 2004 paper do not seem to be the best representation of the real kidney market. As economists have advanced into the area of matching under general constraints and dynamic matching, they have attempted to employ other mechanisms different from TTC. For instance, because all transplantations in any kidney exchange need to be carried out simultaneously, long cycles that could be conducted

⁶How the authors proceed from the setup is notable as they approach the problem from a linear programming perspective. Formulating the matching problem as a linear program and applying the celebrated Scarf’s lemma, they find a random matching that satisfies a notion of stability. They then use an iterative rounding method to find an actual matching (corresponding to a 0–1 solution) such that the resulting matching satisfies stability. Such rounding corresponds to the perturbation of the capacities.

using the TTC mechanism might not be feasible in practice. Roth et al. (2005) provided strategy-proof, constrained-efficient mechanisms of kidney exchange where only pairwise exchanges are permitted. They showed that finding a constrained-efficient matching in their model relates to the cardinality matching problem discussed in the graph theory literature (in addition, the 2005 paper assumed that each patient is indifferent among all kidneys that are compatible to her, based on certain medical evidence). In a 2007 paper, these same authors showed that under certain conditions on kidney supply and demand levels that could normally be expected, full efficiency can be extracted by using exchanges that involve no more than four pairs.

In the papers discussed above, agents and the market itself are static. What if the exchange pool changes over time? Should we conduct exchanges immediately, or if there is no urgency, is it more efficient to wait? These issues are not addressed formally in the aforementioned papers. Ünver (2010) tackles the question of how to conduct barter exchanges in a centralized mechanism when the agent pool evolves over time: he characterizes the efficient two-way and multi-way exchange mechanisms that maximize total exchange surplus. The study of dynamic matching environments has attracted the interest of not only economists but computer scientists and operation research specialists as well. Notable contributions include Anderson et al. (2015) and Akbarpour et al. (2016). There are still many questions left to be addressed, which makes the kidney exchange market one of the great interests among researchers and practitioners today.

School Choice

The third prominent area matching theory is applied to that of school choice and student assignment policy. School choice has become one of the most important and contentious debates in modern education policy. School choice is a policy that allows parents the opportunity to choose the school their child will attend. Traditionally, children are assigned to public schools according to where they live. Wealthy parents already have school choice, because they can

enroll their children in private schools or have the ability to move to a different district entirely. Supporters have argued that school choice helps lower income families by providing them the freedom to send their children to different schools within and across districts. In addition, the increased competition schools face under school choice should incentivize them to increase their quality (for a more extensive survey on empirical and theoretical literature around school choice, see Pathak (2011)). Since it is not possible to assign each student to her top choice school, a central issue in school choice is the design of a student assignment mechanism. One of the first to rigorously and formally tackle this issue with matching theory is Abdulkadiroğlu and Sönmez (2003). The model they propose, which has been regarded as the canonical model, consists of a set of students and schools where:

1. Each student i has a preference relation $\succ_i$ over the schools.
2. Each school c has capacity q_c and priority ordering $\succ_c$ over the students.

One of the reasons we refer to the school's ordering as a priority ordering is because in some school choice programs, orderings are given exogenously (e.g., mandated by law). Pathak (2011) describes a variety of orderings in different districts as follows: "In Boston's school choice plan, for instance, elementary school applicants obtain walk-zone priority if they reside within 1 mile of the school. In other districts, schools construct an ordering of students, as in two-sided problems. In Chicago, for instance, students applying for admissions to selective high schools take an admissions test." (Later, Boston's school choice plan implemented a reform which eliminated the use of walk-zone priority). When evaluating a matching in this setting, two notions are of primary interest: Pareto efficiency and stability. Stability is defined in the standard manner as in section "[Basic Two-Sided Matching Model](#)," while Pareto efficiency only considers students' allocations and does not take into account the schools' priority ordering (as the literature has grown and evolved, generalizations of the notion

of stability have been discussed, which we will examine in section “[Matching with Constraints](#)”). Abdulkadiroğlu and Sönmez (2003) compare three mechanisms: the student-proposing deferred acceptance algorithm, an adaptation of the top trading cycle mechanism (referred to as TTC in this section), and the Boston mechanism.

As the deferred acceptance mechanism is familiar to the reader by now, we describe the other two mechanisms. We start with the Boston mechanism, which, as its name suggests, was in use in school choice programs in the city of Boston (before being replaced by the deferred acceptance algorithm):

- Step 0: Each school orders students by priority block.⁷ Within each block, students are ordered via a lottery system.
- Step 1: In this step, only the first choices of the students are considered. For each school, consider the students who have listed it as their first choice, and assign seats of the school to these students one at a time following their priority order until either there is no seat left or there is no student left who has listed it as his first choice.

In general, for any $t = 1, 2, \dots$

- Step t : In this step, only the t^{th} choices of the students are considered. For each school, consider the students who have listed it as their t^{th} choice, and assign seats of the school to these students one at a time following their priority order until either there is no seat left or there is no student left who has listed it as his t^{th} choice.

In Boston, the Boston mechanism was originally implemented in July 1999 but was abandoned in 2005. One of the central reasons it was abandoned is that it is not strategy proof for

students, i.e., families have an incentive to strategically misreport their preferences. Variations of this mechanism, however, are common in many other school districts.

The next procedure presented by Abdulkadiroğlu and Sönmez (2003) is the TTC mechanism, which is implemented in the following manner (the description of the mechanism is taken directly from Abdulkadiroğlu and Sönmez (2003)):

- Step 1: Assign a counter for each school which keeps track of how many seats are still available at the school. Initially set the counters equal to the capacities of the schools. Each student points to her favorite school under her announced preferences. Each school points to the student who has the highest priority for the school. Since the number of students and schools are finite, there is at least one cycle. Moreover, each school can be part of at most one cycle. Similarly, each student can be part of at most one cycle. Every student in a cycle is assigned a seat at the school she points to and is removed. The counter of each school in a cycle is reduced by one, and if it reduces to zero, the school is also removed. The counters of the schools not in a cycle remain the same.

In general, for any $t = 1, 2, \dots$

- Step t : Each remaining student points to her favorite school among the remaining schools, and each remaining school points to the student with highest priority among the remaining students. There is at least one cycle. Every student in a cycle is assigned a seat at the school that she points to and is removed. The counter of each school in a cycle is reduced by one, and if it reduces to zero, the school is also removed.

This algorithm is very similar to the top trading cycle mechanisms described in sections “House Exchange” and “House Allocation with Existing Owners,” except that agents are not initially endowed with any good. In this adaptation, students are essentially swapping priority orderings with each other. Note that if every school has the same priority ordering, this mechanism reduces to serial

⁷In Boston, first priority consisted of students who lived in a proximal neighborhood and had a sibling that attended the school. The second tier consisted of students with a sibling at the school. Third priority is of the students who live in the “relevant” area. Finally, the remaining students are grouped within the last priority block.

dictatorship where the rank ordering is determined by the priority ranking.

Some of the main properties of TTC are different from those of the deferred acceptance mechanism although both mechanisms are strategy proof. The student-proposing deferred acceptance mechanism is stable, but the resulting outcome is not necessarily Pareto efficient for students, while the top trading cycle mechanism is not stable but produces a Pareto-efficient outcome for students. Whether efficiency or stability is more important is a question that may be an important determinant for the choice of the mechanism. Also, it is then natural to ask whether one can construct an efficient, strategy-proof mechanism which also produces a stable outcome whenever it exists. Kesten (2010) shows that this is impossible.

Much work within school choice literature has expounded on the results of Abdulkadiroğlu and Sönmez (2003). It is important, though, to address criticisms and weakness of the model as well as some difficulties in application of matching theory to the analysis of the policy.

One of the central assumptions of the above model is that students have an exogenous preference over schools that is independent of the other students who are assigned to the same school. This is rather problematic if the quality of a school is affected by the composition of the student body (this is referred to as *peer effect*). The second issue is that the effect of a school choice mechanism on school quality is exogenously given and fixed in the canonical model, although the issue of improving schools takes a center stage of school choice debate in practice (see Hatfield et al. (2016) for an analysis of this topic). Another major difficulty is that the information submitted by students is ordinal and does not necessarily convey information on preference intensities. Abdulkadiroğlu et al. (2011, 2015) and Carroll (2017) analyze this issue theoretically. Agarwal and Somaini's (2016) empirical analysis on strategic reporting in school choice mechanisms highlighted the importance of further study on mechanisms that use the intensity of student preferences. Given these issues (and others we are not discussing here), it is still not completely clear whether the

current notions of stability and Pareto efficiency are the most relevant measures by which to evaluate school choice mechanisms.

Matching with Constraints

We now proceed to a discussion of a relatively new area of research within matching theory and market design application, matching with constraints. This field seeks to study allocations and matching when characteristics and constraints other than the common individual capacity limits are regarded as desirable or required for feasibility. Schools, hospitals, or firms (to use the language of our previous models) may be not only worried about the obvious limit on total individuals they can accept but also about the quantity of types of individuals that are admitted. With the prevalence of affirmative action and the goal of creating a diverse student/employee body, understanding the implementation and impact of such policies is crucial. The desire for diversity ranges beyond just race and gender: in universities, for instance, having students all interested in one or two academic areas is often considered disadvantageous because it may stymie the intellectual growth of its student population.

Abdulkadiroğlu and Sönmez (2003) model a simple affirmative action policy of type-specific quotas and propose mechanisms that satisfy the affirmative action constraints. Under the same type of affirmative action policy, Abdulkadiroğlu (2005) shows that a stable matching can be found using a strategy-proof, student-proposing deferred acceptance algorithm. These papers pushed affirmative action into mainstream matching literature, whereas traditional papers on affirmative action were based on "classical" mechanism design theory.⁸ Kojima (2012) demonstrated various impossibility results that can arise when attempting to implement affirmative action policies in a matching environment. There are situations where affirmative action

⁸The study of employment discrimination began in the second half of the 20th century. The two main theories of discrimination are a theory based on tastes, pioneered by Becker (1957), and a statistical theory, pushed forth by Phelps (1972) and Arrow (1973). Economists such as Glenn Loury and Roland Fryer have further developed the literature around race-based affirmative action.

policies inevitably hurt every minority student under *any stable matching mechanism*. Furthermore, similar impossibility results hold when using TTC. Hafalir et al. (2013) further expound on these phenomena and show that the use of a “quota” versus “reserve” affirmative action system can have significant consequences on the resulting allocation. With minority reserves, schools give higher priority to minority students up to the point that the minorities fill the reserves. They show that the deferred acceptance algorithm with minority reserves is Pareto superior for students to the one with majority quotas.

Kamada and Kojima (2015) advance the idea by looking at matching environments with more general distributional constraints. One example is the Japan Residency Matching Program which imposes regional caps on the numbers of residents so as to limit the concentration of residents in urban areas such as Tokyo. They point out that the mechanisms used in that market and others with constraints suffer from instability and inefficiency. To remedy this problem, they create a modified version of the deferred acceptance algorithm which is strategy proof for students, constrained efficient, and stable in an appropriate sense. Kamada and Kojima (2016, 2017) and Goto et al. (2017) further explore various stability concepts and characterize environments in which stability and other desirable properties such as strategy proofness can be guaranteed.

There are still many issues and problems in the area that are unresolved and worth pursuing. How to address more general types of constraints, especially lower-bound constraints, is still a difficult problem and being actively studied (see Fragiadakis and Troyan (2016) for instance). New mathematical tools from discrete convex analysis have been applied to matching with constraints (Kojima et al. 2016), but the use of such mathematical tools may warrant further investigation.

Conclusion

As indicated throughout this article, matching theory has expanded vastly since the seminal work by Gale and Shapley (1962). Although the theory has advanced considerably, there are many new questions and issues waiting to be explored further.

To begin with, almost all research in the existing literature defines stability under the assumption of complete information, but this is at best a rough approximation of reality. Liu et al. (2014) investigate stability under incomplete information in two-sided matching markets with transfer, while Bikhchandani (2017) studies a similar concept in the no-transfer setting.

Once incomplete information is taken seriously, it is natural to consider “informational externality,” i.e., interdependence in valuations. Chakraborty et al. (2010, 2015) study two-sided matching with interdependent values, while Che et al. (2015) study one-sided matching with interdependent values. In both cases, the possibility of extending desirable matching mechanisms from the standard private value models proved to be severely limited. Designing satisfactory mechanisms under interdependent values is a promising, if challenging, avenue for future research (see Hashimoto (2016) and Pakzad-Hurson (2016) for notable advances).

Another important limitation of the existing literature is that the models tend to be static. Although some matching markets could be approximated well by a static model (e.g., yearly medical residency matching or school choice), others may be better modeled as a dynamic market (e.g., day care slot assignment with arrival and departure of children and the ongoing kidney exchange program). In addition to papers on dynamic kidney exchange already discussed, there is a burgeoning literature on dynamic two-sided matching markets. Kurino (2009), Du and Livne (2016), Doval (2017), and Kadam and Kotowski (2017) propose concepts of dynamic stability and analyze existence under various assumptions on commitment technologies and preferences. This literature is so young that several alternative stability concepts are being studied, but a consensus on the appropriate definition has not been reached yet. In the future, a consensus on the appropriate stability definition may emerge, but it is also possible that different stability concepts are appropriate in different types of dynamic markets. Reaching conclusions on this and other questions awaits further research.

Bibliography

- Abdulkadiroğlu A (2005) College admissions with affirmative action. *Int J Game Theory* 33:535–549
- Abdulkadiroğlu A, Sönmez T (1998) Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66(3):689
- Abdulkadiroğlu A, Sönmez T (1999) House allocation with existing tenants. *J Econ Theory* 88(2):233–260
- Abdulkadiroğlu A, Sönmez T (2003) School choice: a mechanism design approach. *Am Econ Rev* 93(3):729–747
- Abdulkadiroğlu A, Sönmez T (2013) Matching markets: theory and practice. In: Acemoglu D et al (eds) *Advances in economics and econometrics*. Cambridge University Press, Cambridge
- Abdulkadiroğlu A, Che YK, Yasuda Y (2011) Resolving conflicting preferences in school choice: the “Boston mechanism” reconsidered. *Am Econ Rev* 101(1):399–410
- Abdulkadiroğlu A, Che YK, Yasuda Y (2015) Expanding “choice” in school choice. *Am Econ J Microecon* 7(1):1–42
- Agarwal N Somaini (2016) Demand analysis using strategic reports: an application to a school choice mechanism. Working paper
- Akbarpour M et al (2016) Thickness and information in dynamic matching markets. Working paper
- Anderson R, Ashlagi I, Gamarnik D, Kanoria Y (2015) Efficient dynamic barter exchange. *Oper Res* (forthcoming)
- Arrow K (1973) The theory of discrimination. In: Pascal AH (ed) *Racial discrimination in economic life*. D.-C. Heath, Lexington
- Ashlagi I, Braverman M, Hassidim A (2014) Stability in large matching markets with complementarities. *Oper Res* 62(4):713–732
- Azevedo EM Hatfield JW (2017) Existence of equilibrium in large matching markets with complementarities. Working paper
- Becker GS (1957) *The economics of discrimination*. University of Chicago Press, Chicago
- Bikchandani S (2017) Stability with One-sided Incomplete Information. *Journal of Economic Theory* 168:372–399
- Biro P, Klijn F (2013) Matching with couples: a multi-disciplinary survey. *Int Game Theory Rev* 15(2):1–18
- Bogomolnaia A, Moulin H (2001) A new solution to the random assignment problem. *J Econ Theory* 100(2):295–328
- Budish E, Cantillon E (2012) The multi-unit assignment problem: theory and evidence from course allocation at Harvard. *Am Econ Rev* 102(5):2237–2271
- Carroll G (2017) On mechanisms eliciting ordinal preferences. Working paper
- Chakraborty A, Citanna A, Ostrovsky M (2010) Two-sided matching with interdependent values. *J Econ Theory* 145(1):85–105
- Chakraborty A, Citanna A, Ostrovsky M (2015) Group stability in matching with interdependent values. *Rev Econ Des* 19(1):3–24
- Che YK, Kim J, Kojima F (2015) Efficient assignment with interdependent values. *J Econ Theory* 158:54–86
- Che YK, Kim J, Kojima F (2017) Stable matching in large markets. Working paper
- Chen Y, Sönmez T (2006) School choice: an experimental study. *J Econ Theory* 127(1):202–231
- Doval L (2017) A theory of stability in dynamic matching markets. Working paper
- Du S, Livne Y (2016) Rigidity of transfers and unraveling in matching markets. Working paper
- Fragiadakis D Troyan P (2016) Improving matching under hard distributional constraints. *Theor Econ* (forthcoming)
- Gale D, Shapley LS (1962) College admissions and the stability of marriage. *Am Math Mon* 68(1):9–15
- Goto M et al (2017) Designing matching mechanisms under general distributional constraints. *Am Econ J Microecon* (forthcoming)
- Hafalir IE, Yenmez MB, Yildirim MA (2013) Effective affirmative action in school choice. *Theor Econ* 8(2):325–363
- Hashimoto T (2016) The generalized random priority mechanism with budgets. Working paper
- Hassidim A et al (2017) The mechanism is truthful, why aren't you? *American Economic Review Papers and Proceedings* 107:220–224 (forthcoming)
- Hatfield JW, Kojima F (2008) Matching with contracts: comment. *Am Econ Rev* 98:1189–1194
- Hatfield JW, Milgrom PR (2005) Matching with contracts. *Am Econ Rev* 95(4):913–935
- Hatfield JW, Kojima F, Narita Y (2016) Improving schools through school choice: a mechanism design approach. *J Econ Theory* 166:186–211
- Hylland A, Zeckhauser R (1979) The efficient allocation of individuals to positions. *J Polit Econ* 87(2):293–314
- Kadam S Kotowski M (2017) Multi-period matching. Working paper
- Kamada Y, Kojima F (2015) Efficient matching under distributional constraints: theory and applications. *Am Econ Rev* 105(1):67–99
- Kamada Y, Kojima F (2016) Stability and strategy-proofness for matching with constraints: a necessary and sufficient condition. Working paper
- Kamada Y, Kojima F (2017) Stability concepts in matching under distributional constraints. *J Econ Theory* 168:107–142
- Kesten O (2010) School choice with consent. *Q J Econ* 125(3):1297–1348
- Kojima F (2012) School choice: Impossibilities for affirmative action. *Games and Economic Behavior* 75:685–693
- Kojima F (2015) Recent developments in matching theory and its practical applications. *Advances in economics and econometrics*. Cambridge University Press, Cambridge

- Kojima F, Troyan P (2011) Matching and market design: an introduction to selected topics. *Jpn Econ Rev* 62(1):82–98
- Kojima F et al (2013) Matching with couples: stability and incentives in large markets. *Q J Econ* 128(4):1585–1632
- Kojima F, Tamura A, Yokoo M (2016) Designing matching mechanisms under constraints: an approach from discrete convex analysis. Working paper
- Kurino M (2009) Credibility, efficiency, and stability: a theory of dynamic matching markets. Working paper
- Liu Q, Mailath GJ, Postlewaite A, Samuelson L (2014) Stable matching with incomplete information. *Econometrica* 82(2):541–587
- Nguyen T, Vohra R (2017) Near feasible stable matchings with couples. Working paper
- Pakzad-Hurson B (2016) Crowdsourcing and optimal market design. Working Paper
- Pathak PA (2011) The mechanism design approach to student assignment. *Ann Rev Econ* 3(1):513–536
- Pathak PA (2015) What really matters in designing school choice mechanisms. *Advances in economics and econometrics*. Cambridge University Press, Cambridge
- Phelps ES (1972) The statistical theory of racism and sexism. *Am Econ Rev* 62(4):659–661
- Rees-Jones A (2017) Mistaken play in the deferred acceptance algorithm: implications for positive assortative matching. *American Economic Review Papers and Proceedings* 107:225–229 (forthcoming)
- Roth AE (1982) Incentive compatibility in a market with indivisible goods. *Econ Lett* 9(2):127–132
- Roth AE (1984) The evolution of the labor market for medical interns and residents: a case study in game theory. *J Polit Econ* 92(6):991–1016
- Roth AE (1985) The college admissions problem is not equivalent to the marriage problem. *J Econ Theory* 36(2):277–288
- Roth AE (1991) A natural experiment in the organization of entry-level labor markets: regional markets for new physicians and surgeons in the United Kingdom. *Am Econ Rev* 81(3):415–440
- Roth AE (2008a) Deferred acceptance algorithms: history, theory, practice, and open questions. *Int J Game Theory* 36:537–569
- Roth AE (2008b) What we have learned from market design. *Econ J* 118(527):285–310
- Roth AE, Peranson E (1999) The redesign of the matching market for American physicians: some engineering aspects of economic design. *Am Econ Rev* 89(4):748–780
- Roth AE, Postlewaite A (1977) Weak versus strong domination in a market with indivisible goods. *J Math Econ* 4(2):131–137
- Roth AE, Sotomayer MAO (1990) Two-sided matching. Cambridge University Press, Cambridge
- Roth AE, Sönmez T, Ünver MU (2004) Kidney exchange. *Q J Econ* 119(2):457–488
- Roth AE, Sönmez T, Ünver MU (2005) Pairwise kidney exchange. *J Econ Theory* 125(2):151–188
- Roth AE, Sönmez T, Ünver MU (2007) Efficient kidney exchange: coincidence of wants in markets with compatibility-based preferences. *Am Econ Rev* 97(3):828–851
- Shapley L, Scarf H (1974) On cores and indivisibility. *J Math Econ* 1(1):23–37
- Sönmez T, Ünver MU (2009) Matching, allocation, and the exchange of discrete resources. In: Benhabib J et al (eds) *The handbook of social economics*. Elsevier, Amsterdam
- Sönmez T, Ünver MU (2010) Course bidding at business schools. *Int Econ Rev* 51(1):99–123
- Sönmez T, Ünver MU (2011) Market design for kidney exchange. In: Neeman Z et al (eds) *The handbook of market design*. Oxford University Press, Oxford
- Ünver MU (2010) Dynamic kidney exchange. *Rev Econ Stud* 77(1):372–414

Cost Sharing in Production Economies

Maurice Koster
University of Amsterdam, Amsterdam,
The Netherlands

Article Outline

Glossary
Definition of the Subject
Introduction
Cooperative Cost Games
Noncooperative Cost Games
Continuous Cost Sharing Models
Stochastic Cost Sharing Models
Future Directions
Bibliography

Glossary

Core The *core* of a cooperative cost game $\langle N, c \rangle$ is the set of all coalitionally stable vectors of cost shares.

Cost function A *cost function* relates each level of output of a given production technology to the total of minimally necessary units of input to generate it. It is a non-decreasing function, $c : X \rightarrow \mathbb{R}_+$, where X is the (ordered) space of outputs.

Cost sharing problem A cost sharing problem is an ordered pair (q, c) , where $q \in \mathbb{R}_+^N$ is a profile of individual demands of a fixed and finite group of agents $N = \{1, 2, \dots, n\}$ and c is a cost function.

Cost sharing rule A *cost sharing rule* is a mapping that assigns to each cost sharing problem under consideration a vector of nonnegative cost shares.

Demand game It is a strategic game where agents place demands for output strategically.

Demand revelation game It is a strategic game where agents announce their maximal contribution strategically.

Game theory It is the branch of applied mathematics and economics that studies situations where players make decisions in an attempt to maximize their returns. The essential feature is that it provides a formal modeling approach to social situations in which decision-makers interact.

Strategic game An ordered triple $G = \langle N, (A_i)_{i \in N}, (\preceq_i)_{i \in N} \rangle$ where

- $N = \{1, 2, \dots, n\}$ is the set of *players*.
- A_i is the set of available *actions* for player i .
- $\preceq_i$ is a *preference relation* over the set of possible consequences C of action.

Definition of the Subject

Throughout we will use a fixed set of agents $N = \{1, 2, \dots, n\}$ where n is a given natural number. For subsets S, T of N , we write $S \subset T$ if each element of S is contained in T ; $T \setminus S$ denotes the set of agents in T except those in S . The *power set* of N is the set of all subsets of N ; each coalition $S \subset N$ will be identified with the element $\mathbf{1}_S \in \{0, 1\}^N$, the vector with i -th coordinate 1 precisely when $i \in S$. Fix a vector $x \in \mathbb{R}^N$ and $S \subset N$. The projection of x on $\mathbb{R}^S$ is denoted x_S , and $x_{N \setminus S}$ is sometimes more conveniently denoted x_{-S} . For any $y \in \mathbb{R}^S$, (x_{-S}, y) stands for the vector $z \in \mathbb{R}^N$ such that $z_i = x_i$ if $i \in N \setminus S$ and $z_i = y_i$ if $i \in S$. We denote $x(S) = \sum_{i \in S} x_i$. The vector in $\mathbb{R}^S$ with all coordinates equal zero is denoted $\mathbf{0}_S$. Other notation will be introduced when necessary.

This entry focuses on different approaches in the literature through a discussion of a couple of basic and illustrative models, each involving a single facility for the production of a finite set M of outputs, commonly shared by a fixed set $N := \{1, 2, \dots, n\}$ of agents. The feasible set of outputs for the technology is identified with a set $X \subset \mathbb{R}_+^M$. It is assumed that the users of the tech-

nology may freely dispose over any desired quantity or level of the outputs; each agent i has some demand $x_i \in X$ for output. Each profile of demands $x \in X^N$ is associated with its *cost* $c(x)$, i.e., the minimal amount of the idiosyncratic input commodity needed to fulfill the individual demands. This defines the *cost function* $c : X^N \rightarrow \mathbb{R}_+$ for the technology, comprising all the production externalities. A *cost sharing problem* is an ordered pair (x, c) of a demand profile x and a cost function c . The interpretation is that x is produced, and the resulting cost $c(x)$ has to be shared by the collective N . Numerous practical applications fit this general description of a cost sharing problem.

In mathematical terms, a cost sharing problem is equivalent to a *production sharing* problem where output is shared based on the profile of inputs. However, although many concepts are just as meaningful as they are in the cost sharing context, results are not at all easily established using this mathematical duality. In this sense consider (Leroux 2008) as a warning to the reader, showing that the strategic analysis of cost sharing solutions is quite different from surplus sharing solutions. This monograph will center on cost sharing problems. For further reference on production sharing, see Israelsen (1980), Leroux (2004, 2008), and Moulin and Shenker (1992).

Introduction

In many practical situations, managers or policy-makers deal with private or public enterprises with multiple users. A production technology facilitates its users, causing externalities that have to be shared. Applications are numerous, ranging from environmental issues like pollution and fishing grounds to sharing multipurpose reservoirs, road systems, communication networks, and the Internet. The essence in all these examples is that a manager cannot directly influence the behavior of the users but only indirectly by addressing the externalities through some decentralization device. By choosing the right instrument, the manager may help to shape and control the nature of the resulting individual and aggregate behavior. This is what is usually understood as the

mechanism design or implementation paradigm. The state-of-the-art literature shows for a couple of simple but illustrative cost sharing models that one cannot push these principles too far, as there is often a trade-off between the degree of distributive justice and economic efficiency. Then this is what makes choosing “the” right solution an ambiguous task, certainly without a profound understanding of the basic allocation principles. Now first some examples will be discussed.

Example 4.1 The water resource management problem of the Tennessee Valley Authority (TVA) in the 1930s is a classic in the cost sharing literature. It concerns the construction of a dam in a river to create a reservoir, which can be used for different purposes like flood control, hydroelectric power, irrigation, and municipal supply. Each combination of purposes requires a certain dam height, and accompanying construction costs have to be shared by the purposes. Typical for the type of problem is that up to a certain critical height there are economies of scale as marginal costs of extra height are decreasing. Afterward, marginal costs increase due to technological constraints. The problem here is to allocate the construction costs of a specific dam among the relevant purposes. ◁

Example 4.2 Another illustrative cost sharing problem dating back from the early days in the cost sharing literature (Littlechild and Owen 1973; Littlechild and Thompson 1977) deals with landing fee schedules at airports, so-called airport problems. These were often established to cover the costs of building and maintaining the runways. The cost of a runway essentially depends on the size of the largest type of airplane that has to be accommodated – a long runway can be used by smaller types as well. Suppose there are m types of airplanes and that c_i is the cost of constructing a landing strip suitable for type i . Moreover, index the types from small to large so that $0 = c_0 < c_1 < c_2 \dots c_m$. In the above terminology, the technology can be described by $X = \{0, 1, 2, \dots, m\}$, and the cost function $c : X^N \rightarrow \mathbb{R}_+$ is defined by $c(x) = c_k$ where $k = \max \{x_i | i \in N\}$ is the maximal service level required in x . Suppose that in a given year, N_k is

the set landings of type k airplanes; then the set of users of the runway is $N = \cup_k N_k$. The problem is now to apportion the full cost $c(x)$ of the runway to the users in N , where x is the demand vector given by $x_i = \ell$ if $i \in N_\ell$.

Airport problems describe a wide range of cost sharing problems, ranging from sharing the maintenance cost of a ditch system for irrigation projects (Aadland and Kolpin 1998) to sharing the dredging costs in harbors (Bergantino and Coppejans 1997). $\triangleleft$

Example 4.3 A joint project involves a number of activities for which the estimated durations and precedence relations are known. Delay in each of these components affects the period in which the project can be realized. Then a cost sharing problem arises when the joint costs due to the accumulated delay are shared among the individuals causing the delays. See Brânzei et al. (2002). $\triangleleft$

Example 4.4 In many applications the production technology is given by a network $G = (V, E)$ with nodes V and set of costly edges $E \subset V \times V$ and cost function $c : E \rightarrow \mathbb{R}_+$. The demands of the agents are now parts of the infrastructure, i.e., subsets of E . Examples include the sharing the cost of infrastructure for supply of energy and water,

For example, the above airport problem can be modeled as such with

$$\begin{aligned} V &= \{1, 2, \dots, m\} \cup \{\odot\}, \\ E &= \{(\odot, 1), (1, 2), \dots, (m-1, m)\}. \end{aligned}$$

Graphically, the situation is depicted by the line graph in Fig. 1.

Imagine that the runway starts at the special node $\odot$ and that the edges depict the different pieces of runway served to the players. An airplane of type k is situated at node k and needs all


edges toward $\odot$. The edge left to the k -th node is called $e_k = (k-1, k)$, and the corresponding cost is $c(e_k) = c_k - c_{k-1}$. The demand of an airplane at node k is now described by the edges on the path from node k to $\odot$. $\triangleleft$

Example 4.5 In more general network design problems, a link facilitates a *flow*; for instance, in telecommunication it is data flowing through the network; in road systems it is traffic. Henriot and Moulin (1996) discusses a model where a network planner allocates the fixed cost of a network based on the individual demands being flows. Matsubayashi et al. (2005) and Skorin-Kapov and Skorin-Kapov (2005) discuss *congested* telecommunication networks, where the cost of a link depends on the size of the corresponding flow. Then these positive network externalities lead to a concentration of flow and thus to *hub-like* networks. Economies of scale require cooperation of the users, and the problem now is to share the cost of these so-called *hub-like* networks. $\triangleleft$

Example 4.6 As an insurance against the uncertainty of the future net worth of its constituents, firms are often regulated to hold an amount of riskless investments, i.e., its *risk capital*. Given that returns of normal investments are higher, the difference with the riskless investments is considered as a cost. The sum of the risk capitals of each constituent is usually larger than the risk capital of the firm as a whole, and the allocation problem is to apportion this diversification effect observed in risk measurements of financial portfolios. See Denault (2001). $\triangleleft$

Solving Cost Sharing Problems: Cost Sharing Rules

A *vector of cost shares* for the cost sharing problem (x, c) is an element $y \in \mathbb{R}^N$ with the property that $\sum_{i \in N} y_i = c(x)$. This equality is also called



Cost Sharing in Production Economies, Fig. 1 Graphical representation of an airport problem

budget-balancing condition. Central issue addressed in the cost sharing literature is how to determine the appropriate y . The vast majority of the cost sharing literature is devoted to a mechanistic way of sharing joint costs; given a class of cost sharing problems $\mathcal{P}$, a (simple) formula computes the vector of cost shares for each of its elements. This yields a *cost sharing rule* $\mu: \mathcal{P} \rightarrow \mathbb{R}^N$ where $\mu(P)$ is the vector of cost shares for each $P \in \mathcal{P}$.

At this point it should be clear to the reader that many formulas will see to a split of joint cost, and heading for *the* solution to cost sharing problems is therefore an ambiguous task. The least we want from a solution is that it is consistent with some basic principles of fairness or justice and, moreover, that it creates the right incentives. Clearly, the desirability of solution varies with the context in which it is used and so will the sense of appropriateness. Moreover, the different parties involved in the decision-making process will typically hold different opinions; accountants, economists, production managers, regulators, and others all are looking at the same institutional entity from different perspectives. The existing cost sharing literature is about exploring boundaries of what can be thought of desirable features of cost sharing rules. More important than the rules themselves are the properties that each of them is consistent with. Instead of building a theory on single instances of cost sharing problems, the cost sharing literature discusses structural invariance properties over *classes* of problems. Here the main distinction is made on the basis of the topological properties of the technology, whether the cost sharing problem allows for a *discrete* or *continuous* formulation. For each type of models, divisible or indivisible goods, the state-of-the-art cost sharing literature has developed into two main directions, based on the way individual preferences over combinations of cost shares and (levels of) service are treated. Firstly, there is a stream of research in which individual preferences are not explicitly modeled and demands are considered to be *inelastic*. Roughly, it accommodates the large and vast growing *axiomatic* literature (see, e.g., Moulin 2002; Sprumont and Moulin 2007) and the theory on cooperative cost games (Peleg and Sudhölter 2004; Sudhölter 1998; Tijs and Driessen 1986; Young 1985c).


Secondly, there is the literature on cost sharing models where individual preferences are explicitly modeled and demands are elastic. The focus is on noncooperative *demand games* in which the agents are assumed to choose their demands strategically (see, e.g., Kolpin and Wilbur 2005; Moulin and Shenker 1992; Watts 2002).

As an interested reader will soon find out, in the literature there is no shortage of plausible cost sharing techniques. Instead of presenting a kind of summary, this entry focuses on the most basic and most interesting ones and in particular their properties with respect to strategic interplay of the agents.

Outline

The entry is organized as follows: section “[Cooperative Cost Games](#)” discusses cost sharing problems from the perspective of cooperative game theory. Basic concepts like core, Shapley value, nucleolus, and egalitarian solution are treated. Section “[Noncooperative Cost Games](#)” introduces the basic concepts of noncooperative game theory including dominance relations, preferences, and Nash equilibrium. Demand games and demand revelation games are introduced for discrete technologies with concave cost function. This part is concluded with two theorems, the strategic characterization of the Shapley value and constrained egalitarian solution as cost sharing solution, respectively. Section “[Continuous Cost Sharing Models](#)” introduces the continuous production model, and it consists of two parts. First the simple case of a production technology with homogeneous and perfectly divisible private goods is treated. Prevailing cost sharing rules like proportional, serial, and Shapley-Shubik are shortly introduced. We then give a well-known characterization of additive cost sharing rules in terms of corresponding rationing methods and discuss the related cooperative and strategic games. The second part is devoted to the heterogeneous output model and famous solutions like Aumann-Shapley, Shapley-Shubik, and serial rules. In section “[Stochastic Cost Sharing Models](#)” the focus is on a rather new direction in the cost sharing literature, in which the determinants of the cost sharing

**Cost Sharing in
Production Economies,
Fig. 2** Airport game



problem are uncertain and modeled as stochastic variables. Two simple models are discussed in which the *deterministic* cost sharing model is generalized to a more realistic and practical one where both the outcomes of the production technology as well as the costs corresponding to output levels are random variables. The final section “[Future Directions](#)” is looking at future directions of research.

Cooperative Cost Games

A discussion of cost sharing solutions and incentives needs a proper framework wherein the incentives are formalized. In the seminal work of von Neumann and Morgenstern (1944), the notion of a *cooperative game* was introduced as to model the interaction between actors and players who coordinate their strategies in order to maximize joint profits. Shubik (1962) was one of the first to apply this theory in the cost sharing context.

Cooperative Cost Game

A cooperative cost game among players in N is a function $c : 2^N \rightarrow \mathbb{R}$ with the property that $c(\emptyset) = 0$; for non-empty sets $S \subset N$, the value $c(S)$ is interpreted as the cost that would arise should the individuals in S work together and serve only their own purposes. The class of all cooperative cost games for N will be denoted CG .

Any general class $\mathcal{P}$ of cost sharing problems can be embedded in CG as follows. For the cost sharing problem $(x, c) \in \mathcal{P}$ among agents in N , define the *stand-alone* cost game $c_x \in CG$ by

$$c_x(S) := \begin{cases} c(x_S, 0_{N \setminus S}) & \text{if } S \subset N, S \neq \emptyset \\ 0 & \text{if } S = \emptyset. \end{cases} \quad (1)$$

So $c_x(S)$ can be interpreted as the cost of serving only the agents in S .

Example 5.1 The following numerical example will be frequently referred to. An airport is visited by three airplanes in the set $N = \{1, 2, 3\}$ which can be accommodated at cost $c_1 = 12$, $c_2 = 20$, and $c_3 = 33$, respectively. The situation is depicted in Fig. 2.

The corresponding cost game c is determined by associating each coalition S of airplanes to the minimum cost of the runway needed to accommodate each of its members. Then the corresponding cost game c is given by the table below. Slightly abusing notation we denote $c(i)$ to indicate $c(\{i\})$, $c(ij)$ for $c(\{i, j\})$, and so forth.


S	$\emptyset$	1	2	3	12	13	23	123
$c(S)$	0	12	20	33	20	33	33	33

Note that, since we identified coalitions of players in N with elements in 2^N , we may write c to denote the cooperative cost game. By the binary nature of the demands, the cost function for the technology formally is a cooperative cost game. For $x = (1, 0, 1)$ the corresponding cost game c_x is specified by

S	$\emptyset$	1	2	3	12	13	23	123
$c_x(S)$	0	12	0	33	12	33	33	33

Player 2 is a *dummy player* in this game; for all $S \subset N \setminus \{2\}$, it holds $c_x(S) = c_x(S \cup \{2\})$. $\triangleleft$

Example 5.2 Consider the situation as depicted in Fig. 3, where three players, each situated at a different node, want to be connected to the special node $\bullet$ using the indicated costly links. In order to connect themselves to $\bullet$, a coalition S may use only links with $\bullet$ and the direct links between its members and then only if the costs are paid for. For instance, the minimum cost of connecting player 1 in the left node to $\bullet$ is 10, and the cost of connecting players 1 and 2 to $\bullet$ is 18 – the cost



Cost Sharing in Production Economies, Fig. 3 Minimum cost spanning tree problem

of the direct link from 2 and the indirect link between 1 and 2. Then the associated cost game is given by.

S	$\emptyset$	1	2	3	12	13	23	123
$c(S)$	0	10	10	10	18	20	19	27

Notice that in this case the network technology exhibits positive externalities. The more players want to be connected, the lower the per capita cost. $\triangleleft$

For those applications where the cost $c(S)$ can be determined irrespective of the actions taken by its complement $N \setminus S$, the interpretation of c implies *sub-additivity*, i.e., the property that for all $S, T \subset N$ with $S \cap T = \emptyset$ implies $c(S \cup T) \leq c(S) + c(T)$. This is, for instance, an essential feature of the technology underlying natural monopolies (see, e.g., Baumol et al. 1988; Sharkey 1982). Note that the cost games in Examples 5.1 and 5.2 are sub-additive. This is a general property for airport games as well as minimum cost spanning tree games.


Sometimes the benefits of cooperation are even stronger. A game is called *concave* (or *sub-modular*) if for all $S, T \subset N$ we have

$$c(S \cup T) + c(S \cap T) \leq c(S) + c(T). \quad (2)$$

At first this seems a very abstract property, but one may show that it is equivalent with the following, that

$$c(S \cup \{i\}) - c(S) \geq c(T \cup \{i\}) - c(T) \quad (3)$$

for all coalitions $S \subset T \subset N \setminus \{i\}$. This means that the marginal cost of a player i with respect to



Cost Sharing in Production Economies, Fig. 4 Non-concave MCST game

larger coalitions is non-increasing, i.e., the technology exhibits positive externalities. Concave games are also frequently found in the network literature (see Moulin and Shenker 2001; Sharkey 1995; Koster et al. 2002; Maschler et al. 1996).

Example 5.3 Although sub-additive, minimum cost spanning tree games are not always concave. Consider the following example due to Bird (1976). The numbers next to the edges indicate the corresponding cost. We assume a complete graph and that the invisible edges cost 4. Note that in this game, every three-player coalition is connected at cost 12, whereas $c(34) = 16$. Then $c(1234) - c(234) = 16 - 12 = 4$, whereas $c(134) - c(34) = -4$. So the marginal cost of player 1 is not decreasing with respect to larger coalitions (Fig. 4). $\triangleleft$

Incentives in Cooperative Cost Games

The objective in cooperative games is to share the profits or costs savings of cooperation. Similar to the general framework, a *vector of cost shares* for a cost game $c \in CG$ is a vector $x \in \mathbb{R}^N$ such that $x(N) = c(N)$. The question is what cost share vectors make sense if (coalitions of) players have the possibility to opt out thereby destroying cooperation on a larger scale.

In order to ensure that individual players join, a proposed allocation x should at least be *individual rational* so that $x_i \leq c(i)$ for all $i \in N$. In that case no player has a justified claim to reject x as proposal, since going alone yields a higher cost. The set of all such elements is called the *imputation set*. If, in a similar fashion, $x(S) \leq c(S)$ for all

$S \subset N$, then x is called *stable*; under proposal x no coalition S has a strong incentive to go alone, as it is not possible to redistribute the cost shares afterward and make every defector better off. The *core* of a cost game c , notation $\text{core}(c)$, consists of all stable vectors of cost shares for c . If cooperation on a voluntary basis by the grand coalition N is conceived as a desirable feature, then the core and certainly the imputation set impose reasonable conditions for reaching it. Nevertheless, the core of a game can be empty.

Call a collection $\mathcal{B}$ of coalitions *balanced* if there is a vector of positive weights $(\lambda_S)_{S \in \mathcal{B}}$ such that for all $i \in N$

$$\sum_{S \in \mathcal{B}, S \ni i} \lambda_S = 1.$$

A cost game c is *balanced* if it holds for each balanced collection $\mathcal{B}$ of coalition that

$$\sum_{S \in \mathcal{B}} \lambda_S c(S) \geq c(N).$$

It is the celebrated theorem below which characterizes all games with non-empty cores.

Theorem 5.4 *Bondareva-Shapley* (Bondareva 1963; Shapley 1967)

The cost game c is balanced if and only if the core of c is non-empty.

Concave cost games are balanced (see Shapley 1971). Concavity is not a necessary condition for non-emptiness of the core, since minimum cost spanning tree games are balanced as well.

Example 5.5 Consider the two-player game c defined by $c(12) = 10$, $c(1) = 3$, and $c(2) = 8$. Then $\text{core}(c) = \{(x, 10 - x) | 2 \leq x \leq 3\}$. Note that, opposed to the general case, for two-player games, sub-additivity is equivalent with non-emptiness of the core. $\triangleleft$

Cooperative Solutions

A solution on a subclass A of CG is a mapping $\mu : A \rightarrow \mathbb{R}^N$ that assigns to each $c \in A$ a vector of cost shares $\mu(c)$; $\mu_i(c)$ stands for the charge to player i .

The Separable Cost Remaining Benefit Solution

Common practice among civil engineers to allocate costs of multipurpose reservoirs is the following solution. The *separable cost* for each player (read purpose) $i \in N$ is given by $s_i = c(N) - c(N \setminus \{i\})$ and the *remaining benefit* by $r_i = c(i) - s_i$. The separable cost remaining benefit solution charges each player i for the separable cost s_i , and the non-separable costs $c(N) - \sum_{j \in N} s_j$ are then allocated in proportion to the remaining benefits r_i , leading to the formula

$$SCRB_i(c) = s_i + \frac{r_i}{\sum_{j \in N} r_j} \left[c(N) - \sum_{j \in N} s_j \right]. \quad (4)$$

In this formula it is assumed that c is at least sub-additive to ensure that the r_i 's are all positive. For the two-player game c in Example 5.5, the solution is given by $SCRB(c) = (2 + \frac{1}{2}(10 - 9), 7 + \frac{1}{2} \times (10 - 9)) = (2\frac{1}{2}, 7\frac{1}{2})$. In earlier days the solution was well known as “the alternate cost avoided method” or “alternative justifiable expenditure method.” For references, see Young (1985b).

Shapley Value

One of the most popular and oldest solution concepts in the literature on cooperative games is due to Shapley (1953), named *Shapley value*. Roughly it measures the average marginal impact of players. Consider an ordering of the players $\sigma : N \rightarrow N$ so that $\sigma(i)$ indicates the i -th player in the order. Let $\sigma^*(i)$ be the set of the first i players according to σ , so $\sigma^*(1) = \{\sigma(1)\}$, $\sigma^*(2) = \{\sigma(1), \sigma(2)\}$, etc. The *marginal cost share vector* $m^\sigma(c) \in \mathbb{R}^N$ is defined by $m_{\sigma(1)}^\sigma(c) = c(\sigma(1))$, and for $i = 2, 3, \dots, n$

$$m_{\sigma(i)}^\sigma(c) = c(\sigma^*(i)) - c(\sigma^*(i-1)). \quad (5)$$

So according to m^σ , each player is charged with the increase in costs when joining the coalition of players before her. Then the *Shapley value* for c is defined as the average of all $n!$ marginal vectors, i.e.,

$$\Phi(c) = \frac{1}{n!} \sum_{\sigma} m^\sigma(c). \quad (6)$$

Example 5.6 Consider the airport game in Example 5.1. Then the marginal vectors are given by

σ	(123)	(132)	(213)	(231)	(312)	(321)
$\mu^\sigma(c)$	(12,8,13)	(12,0,21)	(0,20,13)	(0,20,13)	(0,0,33)	(0,0,33)

Hence the Shapley value of the corresponding game is $\Phi(c) = (4, 8, 21)$. Following Littlechild and Owen (1973) and Potters and Sudhölter (1999), for airport games this allocation is easily interpreted as the allocation according to which each player pays an equal share of the cost of only those parts of the runway she uses. Then $c(e_1)$ is shared by all three players, $c(e_2)$ only by players 2 and 3, and, finally, $c(e_3)$ is paid in full by player 3. This interpretation extends to the class of standard fixed tree games, where instead of the *lattice* structure of the runway, there is a cost of a tree network to be shared (see Koster et al. 2002). $\triangleleft$

If cost game is concave, then the Shapley value is in the core. Since then each marginal vector specifies a core element and in particular the Shapley value as a convex combination of these. Reconsider the minimum cost spanning tree game c in Example 5.3, a non-concave game with non-empty core and $\Phi(c) = (2\frac{2}{3}, 2\frac{2}{3}, 6\frac{2}{3}, 4)$. Note that this is not a stable cost allocation since the coalition $\{2, 3\}$ would profit by defecting, $c(23) = 8 < 9\frac{1}{3} = \Phi_2(c) + \Phi_3(c)$. Iñarra and Isategui (1993) shows that in general games $\Phi(c) \in \text{core}(c)$ precisely when c is *average concave*. Although not credited as a core selector, the classic way to defend the Shapley value is by the following properties.

Symmetry

Two players i and j are called *symmetric* in the cost game c ; if for all coalitions S not containing i, j , it holds $c(S \cup \{i\}) = c(S \cup \{j\})$. A solution μ is *symmetric* if symmetric players in a cost game c get the same cost shares. If the cost game does not provide any evidence to distinguish between two players, symmetry is the property endorsing equal cost shares.

Dummy

A player i in a cost game c is *dummy* if $c(S \cup \{i\}) = c(S)$ for all coalitions S . A solution μ satisfies *dummy* if $\mu_i(c) = 0$ for all dummy players i in c . So when a player has no impact on costs whatsoever, she cannot be held responsible.

Additivity

A solution is *additive* if for all cost games c_1, c_2 it holds that

$$\mu(c_1) + \mu(c_2) = \mu(c_1 + c_2). \quad (7)$$

For accounting reasons, in multipurpose projects it is a common procedure to subdivide the costs related to the different activities (players) into cost categories, like salaries, maintenance costs, marketing, etc. Each category c_ℓ is associated with a cost game c_ℓ where $c_\ell(S)$ is the total of category ℓ cost made for the different activities in S ; then $c(S) = \sum_\ell c_\ell(S)$ is the joint cost for S . Suppose a solution is applied to each of the cost categories separately; then under an additive solution, the aggregate cost share of an activity is independent from the particular cross section in categories.

Theorem 5.7 Shapley (1953)

Φ is the unique solution on CG which satisfies all three properties *dummy*, *symmetry*, and *additivity*.

Note that *SCRB* satisfies *dummy* and *symmetry* but that it does not satisfy *additivity*. The Shapley value is credited with other virtues, like the following due to Young (1985b).

Consider the practical situation that several division managers simultaneously take steps to increase efficiency by decreasing joint costs, but one division manager establishes a greater relative improvement in the sense that its marginal contribution to the cost associated with all

possible coalitions increases. Then it is more than reasonable that this division should not be penalized. In a broader context, this envisions the idea that each player in the cost game should be credited with the merits of “uniform” technological advances.

Strong Monotonicity

Solution μ is *strongly monotonic* if for any two cost games $c, \bar{c}$ it holds for all $i \in N$ that $c(S \cup \{i\}) - c(S) \geq \bar{c}(S \cup \{i\})$ for all $S \subseteq N \setminus \{i\}$ implies $\mu_i(c) \geq \mu_i(\bar{c})$.

Anonymity is the classic property for solutions declaring independence of solution with respect to the name of the actors in the cost sharing problem. See, e.g., Albizuri and Zarzuelo (2007), Moulin and Shenker (1992), and Pérez-Castrillo and Wettstein (2006). Formally, the definition is as follows. For a given permutation $\pi : N \rightarrow N$ and $c \in CG$, define $\pi c \in CG$ by $\pi c(S) = c(\pi(S))$ for all $S \subseteq N$.

Anonymity

Solution μ is *anonymous* if for all permutations π of N , and all $i \in N$, $\mu_{\pi(i)}(\pi c) = \mu_i(c)$ for all cost games c .

Theorem 5.8 Young (1985b)

The Shapley value is the unique anonymous and strongly monotonic solution.

Myerson (1980) introduced the *balanced contributions axiom* for the model of *nontransferable utility games*, or games without side payments (see Shapley 1969). Within the present context of CG , a solution μ satisfies the *balanced contributions axiom*, if for any cost game c and for any non-empty subset $S \subseteq N$, $\{i, j\} \subseteq S \subseteq N$ it holds that

$$\begin{aligned} \mu_i(S, c) - \mu_i(S \setminus \{j\}, c) \\ = \mu_j(S, c) - \mu_j(S \setminus \{i\}, c). \end{aligned} \quad (8)$$

The underlying idea is the following. Suppose that players agree on using solution μ and that coalition S forms. Then $\mu_i(S, c) - \mu_i(S \setminus \{j\}, c)$ is the amount player i gains or loses when S is already formed and player j resigns. The balanced

contributions axiom states that the gains and/or losses by other player's withdrawal from the coalition should be the same.

Theorem 5.9 Myerson (1980)

There is a unique solution on CG that satisfies the balanced contributions axiom, and that is Φ .

The balanced contribution property can be interpreted in a bargaining context as well. In the game c and with solution μ , a player i can object against player j to the solution $\mu(c)$ when the cost share for j increases when i steps out of the cooperation, i.e., $\mu_j(N, c) \geq \mu_j(N \setminus \{i\})$. In turn, a *counter objection* by player j to this objection is an assertion that player i would suffer more when j ends cooperation, i.e., $\mu_j(N, c) - \mu_j(N \setminus \{i\}) \leq \mu_i(N, c) - \mu_i(N \setminus \{j\})$. The balanced contribution property is equivalent to the requirement that each objection is balanced by a counter objection. For an excellent overview of ideas developed in this spirit, see Maschler (1992).

Another *marginalistic* approach is by Hart and Mas-Colell (1989). Denote for $c \in CG$ the game restricted to the players in $S \subseteq N$ by (S, c) . Given a function $P : CG \rightarrow \mathbb{R}$ which associates a real number $P(N, c)$ to each cost game c with player set N , the *marginal cost* of a player i is defined to be $D^i P(c) = P(N, c) - P(N \setminus \{i\}, c)$. Such a function P with $P(\emptyset, c) = 0$ is called *potential* if $\sum_{i \in N} D^i P(N, c) = c(N)$.

Theorem 5.10 Hart and Mas-Colell (1989)

There exists a unique potential function P , and for every $c \in CG$, the resulting payoff vector $DP(N, c)$ coincides with $\Phi(c)$.

Egalitarian Solution

The Shapley value is one of the first solution concepts proposed within the framework of cooperative cost games, but not the most trivial. This would be to neglect all asymmetries between the players and split total costs equally between them. But as one can expect, egalitarianism in this pure form will not lead to a stable allocation. Just consider the two-player game in Example 5.5 where pure egalitarianism would dictate the

allocation (5, 5), which violates individual rationality for player 1.

In order to avoid these problems, of course we can propose to look for the most egalitarian allocation *within* the core (see Arin and Iñarra 2001; Dutta and Ray 1991). Then in this line of thinking, what is needed in Example 5.5 is a minimal transfer of cost 2 to player 2, leading to the final allocation (3, 7) – the *constrained egalitarian solution*. Although in the former example it was clear what allocation to choose, in general we need a tool to evaluate allocations for the degree of egalitarianism. The earlier mentioned papers all suggest the use of Lorenz order (see, e.g., Atkinson 1970). More precisely, consider two vectors of cost shares x and x' such that $x(N) = x'(N)$. Assume that these vectors are ordered in decreasing order so that $x_1 \geq x_2 \geq \dots \geq x_n$ and $x'_1 \geq x'_2 \geq \dots \geq x'_n$. Then x *Lorenz-dominates* x' – read x is more egalitarian than x' – if for all $k = 1, \dots, n - 1$ it holds that

$$\sum_{i=1}^k x_i \leq \sum_{i=1}^k x'_i, \quad (9)$$

with at least one strict inequality. That is, x is better for those paying the most.

Example 5.11 Consider the three allocations of cost 15 among three players $x = (6, 5, 4)$, $x' = (6, 6, 3)$, and $x'' = (7, 4, 4)$. Firstly, x Lorenz-dominates x'' since $x_1 = 6 < 7 = x''_1$ and $x_1 + x_2 = x'_1 + x'_2$. Secondly, x Lorenz-dominates x' since $x_1 = x'_1$, $x_1 + x_2 < x'_1 + x'_2$. Notice, however, that on the basis of only (9), we cannot make any judgment what is the more egalitarian of the allocations x' and x'' . Since we have $x'_1 = 6 < 7 = x''_1$ but $x'_1 + x'_2 = 6 + 6 > 7 + 4 = x''_1 + x''_2$. The Lorenz order is only a partial order. $\triangleleft$

The constrained egalitarian solution is the set of Lorenz-undominated allocations in the core of a game. Due to the partial nature of the Lorenz order, there may be more than one Lorenz-undominated element in the core. And what if the core is empty? The constrained egalitarian solution is obviously not a straightforward

solution. The original idea of constrained egalitarianism as in Dutta and Ray (1989) focuses on the Lorenz core instead of the core. It is shown that there is at most one such allocation that may exist even when the core of the underlying game is empty.

For concave cost games c , the allocation is well-defined and denoted $\mu^E(c)$. In particular this holds for airport games. Intriguingly, empirical studies (Aadland and Kolpin 1998, 2004) show there is a tradition in using the solution for this type of problems.

For concave cost games c , there exists an algorithm to compute $\mu^E(c)$. This method, due to Dutta and Ray (1989), performs the following consecutive steps. First determine the maximal set S_1 of players minimizing the per capita cost $c(S)/|S|$, where $|S|$ is the size of the coalition S . Then each of these players in S_1 pays $c(S_1)/|S_1|$. In the next step, determine the maximal set S_2 of players in $N \setminus S_1$ minimizing $c_2(S)/|S|$, where c_2 is the cost game defined by $c_2(S) = c(S_1 \cup S) - c(S_1)$. The players in S_2 pay $c_2(S_2)/|S_2|$ each. Continue in this way just as long as not everybody is allocated a cost share. Then in at most n steps, this procedure results in an allocation of total cost, the constrained egalitarian solution. In short the algorithm is as follows:

- **Stage 0:** Initialization, put $S_0^* = \emptyset$, $x^* = 0_N$, and go to stage $t = 1$.
- **Stage t :** Determine

$$S_t \in \arg \max_{S \neq \emptyset} \frac{c(S \cup S_{t-1}^*) - c(S_{t-1}^*)}{|S|}.$$

Put $S_t^* = S_{t-1}^* \cup S_t$ and for $i \in S_t$,

$$x_i^* := \frac{c(S_t^*) - c(S_{t-1}^*)}{|S_t|}.$$

If $S_t^* = N$, we are finished; put $\mu^E(c) = x^*$. Else

repeat the stage with $t := t + 1$.


For example, this algorithm can be used to calculate the constrained egalitarian solution for the airport game in Example 5.1. In the first step

we determine $S_1 = \{1, 2\}$, together with cost shares 10 for the players 1 and 2. Player 3 is allocated the remaining cost in the next step; hence the corresponding final allocation is $\mu^E(c) = (10, 10, 13)$.

Example 5.12 Consider the case where six players share the cost of the following *tree* network that connects them to $\odot$. The *standard fixed tree game* c for this network associates to each coalition of players the minimum cost of connecting each member to $\odot$, where it may use all the links of the tree. This type of games is known to be concave; we can use the above algorithm to calculate $\mu^E(c)$. In the first step, we determine $S_1 = \{1, 3, 4\}$ and each herein pays 8. Then in the second step, the game remains where the edges e_1 , e_3 , and e_4 connecting S_1 have been paid for. Then it easily follows that $S_2 = \{2, 5\}$, so that players 2 and 5 pay 9 each, leaving 10 as cost share for player 6. Thus, we find $\mu^E(c) = (8, 9, 8, 8, 9, 10)$ (Fig. 5). $\triangleleft$

Nucleolus

Given a cost game $c \in CG$ the *excess* of a coalition $S \subset N$ with respect to a vector $x \in \mathbb{R}^N$ is defined as $e(S, x) = x(S) - c(S)$; it measures



Cost Sharing in Production Economies, Fig. 5 Standard fixed tree

dissatisfaction of S under proposal x . Arrange the excesses of all coalitions $S \neq N, \emptyset$ in decreasing order and call the resulting vector $\vartheta(x) \in \mathbb{R}^{2^n-2}$. A vector of cost shares x will be preferred to a vector y , notation $x \succ y$, whenever $\vartheta(x)$ is smaller than $\vartheta(y)$ in the lexicographic order, i.e., there exists i^* such that for $i \leq i^* - 1$, it holds $\vartheta_i(x) = \vartheta_i(y)$ and $\vartheta_{i^*}(x) < \vartheta_{i^*}(y)$. Schmeidler (1969) showed that in the set of individual rational cost sharing vectors, there is a unique element that is maximal with respect to $\succ$, which is called the *nucleolus*. This allocation, denoted by $\nu(c)$, is based on the idea of egalitarianism that the largest complaints of coalitions should consistently be minimized. The concept gained much popularity as a core selector, i.e., it is a one-point solution contained in the core when it is non-empty. This contrasts with the constrained egalitarian solution which might not be well-defined and the Shapley value which may lay outside the core.

Example 5.13 Consider in Example 5.1 the excesses of the different coalitions with respect to the constrained egalitarian solution $\mu^E(c) = (10, 10, 13)$ and the nucleolus $\nu(c) = (6, 7, 20)$:

S	1	2	3	12	13	23
$e(S, \mu^E(c))$	-2	10	-20	0	-10	-10
$e(S, \nu(c))$	-6	-13	-13	-7	-7	-6

Then the ordered excess vectors are

$$\begin{aligned} \vartheta(10, 10, 13) &= (10, 0, -2, -10, -10, -20), \\ \vartheta(6, 7, 20) &= (-6, -6, -7, -7, -13, -13). \end{aligned}$$

Note that indeed $\vartheta(\nu(c)) \succ \vartheta(\mu^E(c))$ since

$$\vartheta_1(6, 7, 20) = -6 < 10 = \vartheta_1(10, 10, 13).$$

The nucleolus of standard fixed tree games may be calculated as a particular *home-down* allocation, as was pointed out by Maschler et al. (1996). $\triangleleft$

For standard fixed tree games and minimum cost spanning tree games, the special structure of

the technology makes it possible to calculate the nucleolus in polynomial time, i.e., with a number of calculations bounded by a multiple of n^2 (see Granot and Huberman 1984). Sometimes one may even express the nucleolus through a nice formula; Legros (1986) showed a class of cost sharing problems for which the nucleolus equals the SCRB solution. But in general calculations are hard and involve solving a linear program with a number of inequalities which is exponential in n . Skorin-Kapov and Skorin-Kapov (2005) suggests to use the nucleolus on the cost game corresponding to hub games.

Instead of the direct comparison of excesses like above, the literature also discusses weighted excesses as to model the asymmetries of justifiable complaints within coalitions. For instance, the per capita *nucleolus* minimizes maximal excesses which are divided by the number of players in the coalition (see Peleg and Sudhölter 2004).

Cost Sharing Rules Induced by Solutions

Most of the above numerical examples deal with cost sharing problems which have a natural and intuitive representation as a cost game. Then basically on such domains of cost sharing problems, there is no difference between cost sharing rules and solutions. It may seem that the cooperative solutions are restricted to this kind of situations. But recall that each cost sharing problem (x, c) is associated its stand-alone cost game $c_x \in CG$, as in (1). Now let μ be a solution on a subclass of $A \subset CG$ and B a class of cost sharing problems (x, c) for which $c_x \in A$. Then a cost sharing rule $\bar{\mu}$ (Fig. 6) is defined on B through

$$\bar{\mu}(x, c) = \mu(c_x). \quad (10)$$

The general idea is illustrated in the diagram on the left. For example, since the Shapley value is defined on the class of all cost games, it defines a cost sharing rule $\bar{\Phi}$ on the class of all cost sharing problems. The cost sharing rule $\bar{\mu}^E$ is defined on the general class of cost sharing problems with corresponding concave cost game. Cost sharing rules derived in this way, *game*

theoretical rules according to Sudhölter (1998), will be most useful below.

Noncooperative Cost Games

Formulating the cost sharing problem through a cooperative cost game assumes inelastic demands of the players. It might well be that for some player, the private merits of service do not outweigh the cost share that is calculated by the planner. She will try to block the payment when no service at no cost is a preferred outcome. Another aspect is that the technology may operate at a sub-optimal level if benefits of delivered services are not taken into account. Below the focus is on a broader framework with elastic demands, which incorporates preferences of a player are defined over *combinations* of service levels and cost shares. The theory of noncooperative games will provide a proper framework in which we can discuss individual aspirations and efficiency of outcomes on a larger scale.


Strategic Demand Games

At the heart of this noncooperative theory is the notion of a *strategic game*, which models an interactive decision-making process among a group of *players* whose decisions may impact the consequences for others. Simultaneously, each player i independently chooses some available action a_i , and the so-realized *action profile* $a = (a_1, a_2, \dots, a_n)$ is associated with some consequence $f(a)$. Below we will have in mind demands or offered contributions as actions, and consequences are combinations of service levels with cost shares.

Preferences over Consequences

Denote by A the set of possible action profiles and C as the set of all consequences of action.

Cost Sharing in Production Economies,
Fig. 6 Induced cost sharing rules



Throughout we will assume that players have preferences over the different consequences of action. Moreover, such preference relation can be expressed by a *utility function* $u_i : C \rightarrow \mathbb{R}$ such that for $z, z' \in C$ it holds $u_i(z) \leq u_i(z')$ if agent i weakly prefers z' to z . Below the set of consequences for agent $i \in N$ will consist of pairs (x, y) where x is the level of service and y a cost share, so that utilities are specified through multi-variable functions, $(x, y) \mapsto u_i(x, y)$.

Preferences over Action Profiles

In turn define for each agent i and all $a \in A$, $U_i(a) = u_i(f(a))$; then U_i assigns to each action profile the utility of its consequence. We will say that the action profile a' is weakly preferred to a by agent i if $U_i(a) \leq U_i(a')$; U_i is called agent i 's utility function *over action profiles*.

Strategic Game and Nash Equilibrium

A strategic game is an ordered triple $G = \langle N, (A_i)_{i \in N}, (U_i)_{i \in N} \rangle$ where

- $N = \{1, 2, \dots, n\}$ is the set of *players*.
- A_i is the set of available *actions* for player i .
- U_i is player i 's utility function over action profiles.

Rational players in a game will choose optimal actions in order to maximize utility. The most commonly used concept in game theory is that of *Nash equilibrium*, a profile of strategies from where unilateral deviation by a single player does not pay. It can be seen as a *steady state* of action in which players hold correct beliefs about the actions taken by others and act rationally. Important assumption here is the level at which the players understand the game; usually it is taken as a starting point that players *know* the complete description of the game, including the action spaces and preferences of others.

Nash Equilibrium (Nash 1950) An action profile a^* in a strategic game $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ is a Nash equilibrium if, for every player i it holds $u_i(a^*) \geq u_i(a_i, a_{-i}^*)$ for every $a_i \in A_i$.

The literature discusses several refinements of this equilibrium concept. One that will play a role in the games below is that of *strong Nash equilibrium* due to Aumann (1959); it is a Nash equilibrium a^* in a strategic game G such that for all $S \subset N$ and action profile a_S , there exists a player $i \in S$ such that $u_i(a_S, a_{N \setminus S}^*) \leq u_i(a^*)$. This means that a strong Nash equilibrium guarantees stability against coordinated deviations, since within the deviating coalition there is at least one agent who does not strictly improve.

Example 6.1 Consider the following two-player strategic game with $N = \{1, 2\}$, $A_1 = \{T, B\}$ and $A_2 = \{L, M, R\}$. Let the utilities be as in the table below

	L	M	R
T	5,4	2,1	3,2
B	4,3	5,2	2,5

Here player 1 chooses a row and player 2 a column. The numbers in the cells summarize the individual utilities corresponding to the action profiles; the first number is the utility of player 1, the second that of player 2. In this game there is a unique Nash equilibrium, which is the action profile (T, L) . $\triangleleft$

Dominance in Strategic Games

In the game $G = \langle N, (A_i)_{i \in N}, (U_i)_{i \in N} \rangle$, the action $a_i \in A_i$ is *weakly dominated* by $a'_i \in A_i$ if $U_i(a_i, a_{-i}) \leq U_i(a'_i, a_{-i})$ for all $a_{-i} \in A_{-i}$, with strict inequality for some profile of actions a_{-i} . If strict inequality holds for all a_{-i} , then a_i is strictly dominated by $\tilde{a}_i$. Rational players will not use strictly dominated strategies, and, as far as prediction of play is concerned, these may be eliminated from the set of possible actions. If we do this elimination step for each player, then we may reconsider whether some actions are dominated within the reduced set of action profiles. This step-by-step reduction of action sets is called the procedure of *successive elimination of (strictly) dominated strategies*. The set of all action profiles surviving this procedure is denoted D^∞ .

Example 6.2 In Example 6.1 action M of player 2 is strictly dominated by L and R . Player 1 has no dominated actions. Now eliminate M from the actions for player 2. Then the reduced game is

	L	R
T	5,4	3,2
B	4,3	2,5

Notice that action L for player 1 was not dominated in the original game, for the reason that B was the better of the two actions against M . But if M is never played, T is strictly better than B . Now eliminate B , yielding the reduced game

	L	R
T	5,4	3,2

In this game, L dominates R ; hence the only action profile surviving the successive elimination of strictly dominated strategies is (T, L) . $\triangleleft$

A stronger notion than dominance is the following. Call an action $a_i \in A_i$ *overwhelmed* by $a'_i \in A_i$ if

$$\max \{U_i(a_i, a_{-i}) | a_{-i} \in A_{-i}\} < \min \{U_i(a'_i, a_{-i}) | a_{-i} \in A_{-i}\}.$$

Then O^∞ is the set of all action profiles surviving the successive elimination of overwhelmed actions. This notion is due to Friedman and Shenker (1998) and Friedman (2002). In Example 6.1 the action M is overwhelmed by L , not by R . Moreover, the remaining actions in O^∞ are B, T, L , and R .

Demand Games

Strategic games in cost sharing problems arise when we assume that the users of the production technology choose their demands strategically and a cost sharing rule sees to an allocation of the corresponding service costs. The action profiles A_i are simply specified by the demand spaces of the agents, and utilities are specified over combinations of (level of) received service and accompanying cost shares. Hence utilities are defined over consequences of action, and $u_i(q_i, x_i)$ denotes

i 's utility at receiving service level q_i and cost share x_i ; u_i is increasing in the level of received service x_i and decreasing in the allocated cost y_i . Now assume a cost function c and a cost sharing rule μ . Then given a demand profile $a = (a_1, a_2, \dots, a_n)$, the cost sharing rule determines a vector of cost shares $\mu(a, c)$ and in return also the corresponding utilities over demands $U_i(a) = u_i(a_i, \mu_i(a, c))$. Observe that agents influence each other's utility via the cost component. The *demand game* for this situation is then the strategic game

$$G(\mu, c) = \langle N, (A_i)_{i \in N}, (U_i)_{i \in N} \rangle. \quad (11)$$

Example 6.3 Consider the airport problem in Example 5.1. Each player now may request service (1) or not (0). Then the cost function is fully described by the demand of the largest player. That is, $c(x) = 33$ if 3 requires service, $c(x) = 20$ for all profiles with $x_2 = 1, x_3 = 0$, and $c(x) = 12$ if $x = (1, 0, 0)$, $c(0, 0, 0) = 0$. Define the cost sharing rule $\overline{\Phi}(x, c) = \Phi(c_x)$, that is, $\overline{\Phi}$ calculates the Shapley value for the underlying cost game c_x as in (1). Assume that the players' preferences over ordered pairs of service level and cost shares are fully described by

$$\begin{aligned} u_1(q_1, x_1) &= 8q_1 - x_1, \\ u_2(q_2, x_2) &= 6q_2 - x_2, \\ u_3(q_3, x_3) &= 30q_3 - x_3. \end{aligned}$$

Here q_i takes values 0 (no service) or 1 (service) and x_i stands for the allocated cost. So player 1 prefers to be served at unit cost instead of not being served at zero cost, $u_1(0, 0) = 0 < 7 = u_1(1, 1)$. The infrastructure is seen as an excludable public good, so those with demand 0 do not get access to the technology. Each player now actively chooses to be served or not, so her action set is specified by $A_i = \{0, 1\}$. Recall the definition of c_x as in (1). Then given a profile of such actions $a = (a_1, a_2, a_3)$ and cost shares $\overline{\Phi}(a, c)$, utilities of the players in terms of action profiles become $U_i(a) = u_i(a_i, \Phi_i(a, c))$, so that

$$\begin{aligned} U_1(a) &= 8a_1 - \overline{\Phi}_1(a, c), \\ U_2(a) &= 6a_1 - \overline{\Phi}_2(a, c), \\ U_3(a) &= 30a_3 - \overline{\Phi}_3(a, c). \end{aligned}$$

Now that we provided all details of the demand game $G(\overline{\Phi}, c)$, let us look for (strong) Nash equilibria. Suppose that the action profile $a^* = (1, 0, 1)$ is played in the game. Then in turn the complete infrastructure is realized just for players 1 and 3, and the cost allocation is given by $(6, 0, 27)$. Then the vectors of individual utilities are given by $(2, 0, 3)$. Now if we consider unilateral deviations from a^* , what happens to the individual utilities?

$$\begin{aligned} U_1(0,0,1) &= 0 < 2 = U_1(1,0,1), \\ U_2(1,1,1) &= 6 - \overline{\Phi}_2((1,1,1), c) = 6 - \Phi(c) \\ &= 6 - 8 = -2 < 0 = U_2(1,0,1), \\ U_3(1,0,0) &= 0 < 3 = U_3(1,0,1). \end{aligned}$$

This means that for each player, unilateral deviation does not pay; a^* is a Nash equilibrium. The first inequality shows as well why the action profile $(0, 0, 1)$ is not. It is easy to see that the other Nash equilibrium of this game is the action profile $(0, 0, 0)$; no player can afford the completion of the infrastructure just for herself. Notice however that this zero profile is not a strong Nash equilibrium as players 1 and 3 may well do better by choosing for service at the same time, ending up in $(1, 0, 1)$. The latter profile is the unique strong Nash equilibrium of the game.

Similar considerations in the demand game $G(\overline{\mu}^E, c)$ induced by the constrained egalitarian solution lead to the unique strong Nash equilibrium $(0, 0, 0)$; nobody wants service. $\triangleleft$

With cost sharing rules as decentralization tools, the literature postulates Nash equilibria of the related demand game as the resulting behavioral mode. This is a delicate step because – as the example above shows – it is easy to find games with many equilibria, which causes a selection problem. And what can we do if there are no equilibria? This will not be the topic of this text, and the interested reader is referred to any standard textbook on game theory, for instance, see Osborne and Rubinstein (1994), Osborne (2004), and Ritzberger (2002). If there is a unique equilibrium, then it is taken as the prediction of actual play.

Demand Revelation Games

For a social planner, one way to retrieve the level at which to operate the production facility is via a pre-specified demand game. Another way is to ask each of the agents for the maximal amount that she is willing to contribute in order to get service, and then, contingent on the reported amounts, install an adequate level of service together with a suitable allocation of costs. Opposed to demand games ensuring the level of service, in a demand revelation game, each player is able to ensure a maximum charge for service.

The approach will be discussed under the assumption of a discrete production technology with binary demands, so that the cost function c for the technology is basically the characteristic function of a cooperative game. Moreover assume that the utilities of the agents in N are quasi-linear and given by

$$u_i(q_i, x_i) = \alpha_i q_i - x_i \quad (12)$$

where $q_i \in \{0, 1\}$ denotes the service level, x_i stands for the cost share, and α_i is a nonnegative real number. Moulin and Shenker (2001) discusses this framework and assumes that c is concave, and Koster et al. (2003) and Young (1998) moreover take c as the joint cost function for the realization of several discrete public goods.

Demand Revelation Mechanisms

Formally, a *revelation mechanism* M assigns to each profile η of reported maximal contributions a set $S(\eta)$ of agents receiving service and $x(\eta)$ a vector of monetary compensations. Here we will require that these monetary compensations are cost shares; given some cost sharing rule μ , the vector $x(\eta)$ is given by $\mu(\mathbf{1}_{S(\eta)}, c)$ where c is the relevant cost function.

Moreover, note that by restricting ourselves to cost share vectors, we implicitly assume non-positive monetary transfers. The *budget balance* condition is crucial here; otherwise mechanisms of a different nature must be considered as well (see Clarke 1971; Green and Laffont 1977; Groves 1973). There are other ways that a planner may use to determine a suitable service level is by demanding prepayment from the players and

determine a suitable service level on the basis of these labeled contributions (see Koster et al. 2003; Young 1998).

Many mechanisms come to mind, but in order to avoid too much arbitrariness from the planner's side, the more sensible ones will grant the players some control over the outcomes. We postulate the following properties:

- **Voluntary participation (VP).** Each agent i can guarantee herself the welfare level $u_i(0, 0)$ (no service, no payment) by reporting truthfully the maximal willingness to pay, which is α_i under (12).
- **Consumer sovereignty (CS).** For each agent i , a report y_i exists so that she receives service, irrespective of the reports by others.

Now suppose that the planner receives the message $\eta = \alpha$, known to her as the profile of true player characteristics. Then for economic reasons she could choose to serve a coalition S of players that maximizes the *net benefit* at α , $\pi(S, \alpha) = \alpha(S) - c(S)$. However, problems will arise when some player i is supposed to pay more than α_i , so the planner should be more careful than that. She may want to choose a coalition S with maximal $\pi(S, \alpha)$ such that $\mu(\mathbf{1}_S, c) \leq \alpha$ holds; such set S is called *efficient*. But in general the planner cannot tell whether the players reported truthfully or not; what should she do then? One option is that she applies the above procedure thereby naively holding each reported profile η for the *true* player characteristics. In other words, she will pick a coalition that solves the following optimization problem

$$\begin{aligned} \max_{S \subseteq N} \pi(S, \eta) &= \eta(S) - c(S) \\ \text{s.t. } \mu(\mathbf{1}_S, c) &\leq \eta \end{aligned} \quad (13)$$

Denote such a set by $S(\mu, \eta)$. If this set is unique, then the demand revelation mechanism $M(\mu)$ selects $S(\mu, \eta)$ to be served at cost shares determined by $x(\eta) = \mu(\mathbf{1}_S, c)$. This procedure will be explained through some numerical examples.

Example 6.4 Consider the airport problem and utilities of players over service levels and cost

shares as in Example 5.1. Moreover assume the planner uses the Shapley cost sharing rule $\bar{\Phi}$ as in Example 6.3 and that she receives the true profile of preferences from the players, $\alpha = (8, 6, 30)$. Calculate for each coalition S the net benefits at α :

S	$\emptyset$	1	2	3	12	13	23	N
$\pi(S, \alpha)$	0	-4	-14	-3	-6	8	3	11

Not surprisingly, the net benefits are the highest for the grand coalition. But if N were selected by the mechanism, the corresponding cost shares are given by $\bar{\Phi}(\mathbf{1}_N, c) = (4, 8, 21)$, and player 2 is supposed to contribute more than she is willing to. Then the second highest net benefits are generated by serving $S = \{1, 3\}$, with cost shares $\bar{\Phi}(\mathbf{1}_S, c) = (6, 0, 27)$. Then $\{1, 3\}$ is the solution to (13).

What happens if some of the players misrepresent their preferences, for instance, like in $\eta = (13, 6, 20)$? The planner determines the *conceived* net benefits

S	$\emptyset$	1	2	3	12	13	23	N
$\pi(S, \eta)$	0	1	-14	-13	-1	0	-7	6

Again, if the planner served the coalition with the highest net benefit, N , then player 2 would refuse to pay. Second highest net benefit corresponds to the singleton $S = \{1\}$, and this player will get service under $M(\bar{\Phi})$ since $\eta_3 = 13 > 12 = c(1, 0, 0)$. $\triangleleft$

Example 6.5 Consider the same situation but now with $\bar{\mu}^E$ instead of $\bar{\Phi}$ as cost sharing rule. Now consider the following table, restricted to coalitions with nonnegative net benefits (all other will not be selected):

S	$\emptyset$	13	23	N
$\pi(S, \alpha)$	0	8	3	11
$\bar{\mu}^E(\mathbf{1}_S, c)$	(0,0,0)	(12,0,21)	(0, $\frac{33}{2}$, $\frac{33}{2}$)	(10, 10, 13)

Here only the empty coalition $S = \emptyset$ satisfies the requirement $\bar{\mu}^E(\mathbf{1}_S, c) = (0, 0, 0) \leq (8, 6, 30)$;

hence according to the mechanism $M(\bar{\mu}^E)$, nobody will get service. $\triangleleft$

In general, the optimization problem (13) does not give unique solutions in case of which a planner should still further specify what she does in those cases. For concave cost functions, consider the following sequence of coalitions

$$S_1 = N, \quad S_t = \{i \in S_{t-1} \mid \eta_i \geq \mu_i(\mathbf{1}_{S_t}, c)\}.$$

So, starting with the grand coalition N , at each consecutive step, those players are removed whose maximal contributions are not consistent with the proposed cost share – until the process settles down. The set of remaining players defines a solution to (13) and taken to define $S(\mu, \eta)$.

Strategy-Proofness

The essence of a demand revelation mechanism M is that its rules are set up in such a way that it provides enough incentives for the players not to lie about their true preferences. We will now discuss the most common non-manipulability properties of revelation mechanisms in the literature.

Fix two profiles $\alpha', \alpha \in \mathbb{R}_+^N$, where α corresponds to the true maximal willingness to pay. Let (q', x') and (q, x) be the allocations implemented by the mechanism M on receiving the messages α' and α , respectively. The mechanism M is called *strategy-proof* if it holds for all $i \in N$ that $\alpha'_{N \setminus \{i\}} = \alpha_{N \setminus \{i\}}$ implies $u_i(q'_i, x'_i) \leq u_i(q_i, x_i)$.

So, given the situation that the other agents report truthfully, unilateral deviation by agent i from the true preference never produces better outcomes for her. Similarly, M is *group strategy-proof* if deviations of groups of agents does not pay for all deviators, i.e., for all $T \subset N$ the fact that $\alpha'_{N \setminus T} = \alpha_{N \setminus T}$ implies $u_i(q'_i, x'_i) \leq u_i(q_i, x_i)$ for all $i \in T$. So, under a (group) strategy-proof mechanism, there is no incentive to act untruthfully by misrepresenting the true preferences, and this gives a benevolent planner control over the outcome.

Cross-Monotonicity

Cost sharing rule μ is called *cross-monotonic* if the cost share of an agent i is not increasing if other agents demand more service, in case of a concave cost function c . Formally, if $\bar{x} \geq x$ and $\bar{x}_i \geq x_i$, then $\mu_i(\bar{x}, c) \leq \mu_i(x, c)$; each agent is granted a (fair) share of the positive externality due to an increase in demand by others.

Proposition 6.6 Moulin and Shenker (2001)

The only group strategy-proof mechanisms $M(\mu)$ satisfying VP and CS are those related to cross-monotonic cost sharing rules μ .

There are many different cross-monotonic cost sharing rules, and thus just as many mechanisms that are group strategy-proof. Examples include the mechanisms $M(\Phi)$ and $M(\bar{\mu}^E)$, because Φ and $\bar{\mu}^E$ are cross-monotonic. However, the nucleolus is not cross-monotonic and does therefore not induce a strategy-proof mechanism.

Above we discussed two instruments a social planner may invoke to implement a desirable outcome without knowing the true preferences of the agents. Basically, the demand revelation games define a so-called direct mechanism. The announcement of the maximal price for service pins down the complete preferences of an agent. So in fact the planner decides upon the service level based on a complete profile of preferences. In case of a cross-monotonic cost sharing rule, under the induced mechanism truth-telling is a weakly dominant strategy; announcing the true maximal willingness to pay is optimal for the agent regardless of the actions of others. This means good news for the social planner as the mechanism is self-organizing: the agents need not form any conjecture about the behavior of others in order to know what to do. In the literature (Dasgupta et al. 1979), such a mechanism is called *straightforward*. The demand games define an *indirect mechanism*, as by reporting a demand the agents do no more than signaling their preferences to the planner.

Although in general there is a clear distinction between direct and indirect mechanisms, in the model presented in this section, these are nevertheless strongly connected. Focus on a demand game $G(\mu, c)$; recall that in this game, the agents

simultaneously and independently decide upon requesting service or not, and costs are shared using the rule μ among those agents receiving service. Suppose that for each profile of utility functions as in (12) the resulting game $G(\mu, c)$ has a unique (strong) Nash equilibrium. Then this equilibrium can be taken to define a mechanism. That is, the mechanism elicits u and chooses the unique equilibrium *outcome* of the reported demand game. Then this mechanism is equivalent with the demand revelation mechanism. Observe that indeed the strong equilibrium $(1, 0, 1)$ in the game $G(\bar{\Phi}, c)$ in Example 6.3 corresponds to players chosen by $M(\bar{\Phi})$ under truthful reporting. And where no player is served in the strong equilibrium of $G(\bar{\mu}^E, c)$, none of the players is selected by $M(\bar{\mu}^E)$. It is a general result in implementation theory due to Dasgupta et al. (1979) that a direct mechanism constructed in this way is (group) strategy-proof *provided* the underlying space of preferences is *rich*. Bochet and Klaus (2007) shows that for the general result, *richness* in the sense of Maskin and Sjöström (2002) is needed opposed to the version in Dasgupta et al. (1979). It is easily seen that the above sets of preferences meet the requirements. To stress importance of such a structural property as richness, it is instructive to point at what is yet to come in section “Uniqueness of Nash Equilibria in P¹-Demand Games.” Here, the strategic analysis of the demand game induced by the proportional rule shows uniqueness of Nash equilibrium on the domain of preferences $\mathcal{L}^*$ if only costs are convex. However, this domain is *not* rich, and the direct mechanism defined in the same fashion as above by the Nash equilibrium selection is *not* strategy-proof.

Efficiency and Strategy-Proof Cost Sharing Mechanisms

Suppose *cardinal utility* for each agent, so that intercomparison of utility is allowed. Proceeding on the net benefit of a coalition, we may define its *value* at α by

$$v(N, \alpha) = \max_{S \subset N} \pi(S, \alpha), \quad (14)$$

where $\pi(S, \alpha)$ is the net benefit of S at α . A coalition S such that $v(N, \alpha) = \pi(S, \alpha)$ is called *efficient*. It will be clear that a mechanism $M(\mu)$

that is defined through the optimization problem (13) will not implement an efficient coalition of served players, due to the extra constraint on the cost shares.

For instance, in Example 5.1, the value of the grand coalition at $\alpha = (8, 6, 30)$ is given by $v(N, \alpha) = \alpha(N) - c(N) = 44 - 33 = 11$. At the same profile, the implemented outcome by mechanism $M(\bar{\Phi})$ gives rise to a total surplus of $38 - 30 = 8$ for the grand coalition – which is not optimal. The mechanism $M(\bar{\mu}^E)$ performs even worse as it leads to the stand-alone surplus 0; none is served.

This observation holds for far more general settings, and, moreover, it is a well-known result from implementation theory that – under non-constant marginal cost – any strategy-proof mechanism based on full coverage of total costs will not always implement efficient outcomes. For the constant marginal cost case, see Leroux (2004) and Maniquet and Sprumont (1999). Then, if there is an unavoidable loss in using demand revelation mechanisms, can we still tell which mechanisms are more efficient? Is it a coincidence that in the above examples the Shapley value performs better than the egalitarian solution?

The *welfare loss* due to $M(\mu)$ at a profile of true preferences α is given by

$$L(\mu, \alpha) = v(N, \alpha) - \{\alpha(S(\mu, \alpha)) - c(S)\}. \quad (15)$$

For instance, with $\alpha = (8, 6, 30)$ in the above examples, we calculate $L(\bar{\Phi}, \alpha) = 11 - 8 = 3$ and $L(\bar{\mu}^E, \alpha) = 11 - 0 = 11$. An overall measure of quality of a cost sharing rule μ in terms of efficiency loss is defined by

$$\gamma(\mu) = \sup_{\alpha} L(\mu, \alpha). \quad (16)$$

Theorem 6.7 Moulin and Shenker (2001)

Among all mechanisms $M(\mu)$ derived from cross-monotonic cost sharing rules μ , the Shapley rule $\bar{\Phi}$ has the unique smallest maximal efficiency loss, or $\gamma(\bar{\Phi}) < \gamma(\mu)$ if $\mu \neq \bar{\Phi}$.

Notice that this makes a strong case for the Shapley value against the egalitarian solution. The story does however not end here.

Mutuswami (2004) considers a model where the valuations of the agents for the good are independent random variables, drawn from a distribution function F satisfying the *monotone hazard condition*. This means that the function defined by $h(x) = f(x)/(1 - F(x))$ is non-decreasing, where F has f as density function. It is shown that the constrained egalitarian solution maximizes the probability that all members of any given coalition accept the cost shares imputed to them. Moreover, Mutuswami (2004) characterized the solution in terms of efficiency. Suppose for the moment that the planner calculated cost share vector x for the coalition S and that its members are served conditional on acceptance of the proposed cost shares. The probability that all members of S accept the shares is given by $P(x) = \prod_{i \in S} (1 - F(x_i))$, and if we assume that the support of F is $(0, m)$, then the expected surplus from such an offer can be calculated as follows:

$$W(x) = P(x) \cdot \left[\sum_{i \in S} \int_{x_i}^m \frac{u_i}{1 - F(x_i)} dF(u_i) - c(S) \right]. \quad (17)$$

The finding of Mutuswami (2004) is that for *log-concave* f , i.e., $x \mapsto \ln(f(x))$ is concave (An 1998), the mechanism based on the constrained egalitarian solution not only maximizes the probability that a coalition accepts the proposal, but it maximizes its expected surplus as well. Formally,

Theorem 6.8 (Mutuswami 2004)

If the profile of valuations $(u_i)_{i \in N}$ are independently drawn from a common distribution function F with log-concave and differentiable density function f , then $W(\bar{\mu}^E(\mathbf{1}_S, c)) \geq W(\mu(\mathbf{1}_S, c))$ for all cross-monotonic solutions μ and all $S \subset N$.

Extension of the Model: Discrete Goods

Suppose the agents consume idiosyncratic goods produced in indivisible units. Then

given a profile of demands, the cost associated with the joint production must be shared by the users. Then this model generalizes the binary good model discusses so far, and it is a launch pad to the continuous framework in the next section. In this discrete good setting, Moulin (1999) characterizes the cost sharing rules which induce strategy-proof social choice functions defined by the equilibria of the corresponding demand game. As it turns out, these rules are *basically* the *sequential stand-alone rules* according to which costs are shared in an incremental fashion with respect to a fixed ordering of the agents. This means that such a rule charges the first agent for her stand-alone costs, the second for the stand-alone cost for the first two users minus the stand-alone cost of the first, etc. Here the word “basically” refers to all discrete cost sharing problems other than those with binary goods. Then here the sufficient condition for strategy-proofness is that the underlying cost sharing rule be cross-monotonic, which admits other rules than the sequential ones – like $\bar{\mu}^E$ and $\bar{\Phi}$.

Continuous Cost Sharing Models

Continuous Homogeneous Output Model, $\mathcal{P}^1$

This model deals with production technologies for one single perfectly divisible output commodity. Moreover, we will restrict ourselves to private goods. Many ideas below have been studied for *public goods* as well; for further references, see, e.g., Fleurbaey and Sprumont (2009), Maniquet and Sprumont (2004), and Moulin (1994).

The demand space of an individual is given by $X = \mathbb{R}_+$. The technology is described by a non-decreasing cost function $c : \mathbb{R}_+ \rightarrow \mathbb{R}$ such that $c(0) = 0$, i.e., there are no fixed costs. Given a profile of demands $x \in \mathbb{R}_+^N$, costs $c(x(N))$ have to be shared. Moreover, the space of cost functions will be restricted to those c being *absolutely continuous*. Examples include the differentiable and Lipschitz continuous functions. Absolute continuity implies that aggregate costs for production can be calculated by the total of marginal costs

$$c(y) = \int_0^y c'(t) dt.$$

Denote the set of all such cost functions by C^1 and the related cost sharing problems by $\mathcal{P}^1$. Several cost sharing rules on $\mathcal{P}^1$ have been proposed in the literature.

Average Cost Sharing Rule

This is the most popular and oldest concept in the literature and advocates Aristotle's principle of proportionality.

$$\mu^{\text{AV}}(x, c) = \begin{cases} \frac{x}{x(N)} \cdot c(x(N)) & \text{if } x \neq 0_N \\ 0 & \text{if } x = 0_N \end{cases} \quad (18)$$

Shapley-Shubik Rule

Each cost sharing problem $(x, c) \in \mathcal{P}^1$ is related to the stand-alone cost game c_x such that $c_x(S) = c(x(S))$. Then the Shapley-Shubik rule is determined by application of the Shapley value to this game:

$$\mu^{\text{SS}}(x, c) = \Phi(c_x).$$

Serial Rule

This rule, due to Moulin and Shenker (1992), determines cost shares by considering particular intermediate cost levels. More precisely, given $(x, c) \in \mathcal{P}^1$, it first relabels

the agents by increasing demands such that $x_1 \leq x_2 \leq \dots \leq x_n$. The intermediate production levels are

$$y^1 = nx_1, \quad y^2 = x_1 + (n-1)x_2, \dots, \quad y^k = \sum_{j=1}^{k-1} x_j + (n-k+1)x_k, \dots, \quad y^n = x(N).$$



These levels are chosen such that at each new level, one agent more is fully served (Fig. 7).

His demand: at y^1 each agent is handed out x_1 , at y^2 agent 1 is given x_1 and the rest x_2 , etc. The serial cost shares are now given by

$$\mu_i^{\text{SR}}(x, c) = \sum_{\ell=1}^k \frac{c(y^\ell) - c(y^{\ell-1})}{n - \ell + 1}.$$

So according to μ^{SR} , each agent pays a fair share of the incremental costs in each stage that she gets new units.

Example 7.1 Consider the cost sharing problem (x, c) with $x = (10, 20, 30)$ and $c(y) = \frac{1}{2}y^2$. Then first calculate the intermediate production levels



Cost Sharing in Production Economies, Fig. 7 Intermediate production levels

$y^0 = 0$, $y^1 = 30$, $y^2 = 50$, and $y^3 = 60$. Then the cost shares are calculated as follows:

$$\begin{aligned}\mu_1^{\text{SR}}(x, c) &= \frac{c(y^1) - c(y^0)}{3} = 150, \\ \mu_2^{\text{SR}}(x, c) &= \mu_1^{\text{SR}}(x, c) + \frac{c(y^2) - c(y^1)}{2} \\ &= 150 + \frac{1250 - 450}{2} = 550, \\ \mu_3^{\text{SR}}(x, c) &= \mu_2^{\text{SR}}(x, c) + c(y^3) - c(y^2) \\ &= 550 + 550 = 1100.\end{aligned}$$

◁

The serial rule has attracted much attention lately in the network literature and found its way in fair queueing packet scheduling algorithms in routers (Demers et al. 1990).

Decreasing Serial Rule

de Frutos (1998) proposes serial cost shares where demands of agents are put in decreasing order. Resulting is the *decreasing serial rule*. Consider a demand vector $x \in \mathbb{R}_+^N$ such that $x_1 \leq x_2 \leq \dots \leq x_n$. Define recursively the numbers y^ℓ for $\ell = 1, 2, \dots, n$ by $y^\ell = \ell x_\ell + x_{\ell+1} + \dots + x_n$, and put $y^{n+1} = 0$. Then the decreasing serial rule is defined by

$$\mu_i^{\text{DSR}}(x, c) = \sum_{\ell=i}^n \frac{c(y^\ell) - c(y^{\ell+1})}{\ell}. \quad (19)$$

Example 7.2 For the cost sharing problem in Example 7.1, calculate $y^1 = 90$, $y^2 = 70$, $y^3 = 60$, and then

$$\begin{aligned}\mu_3^{\text{DSR}}(x, c) &= \frac{c(y^3) - c(y^4)}{3} = \frac{4050 - 0}{3} = 1350, \\ \mu_2^{\text{DSR}}(x, c) &= \mu_3^{\text{DSR}}(x, c) + \frac{c(y^2) - c(y^3)}{2} \\ &= 1350 + \frac{2450 - 4050}{2} = 550, \\ \mu_1^{\text{DSR}}(x, c) &= \mu_2^{\text{DSR}}(x, c) + (c(y^1) - c(y^2)) \\ &= 550 + (1800 - 2450) = -100.\end{aligned}$$

◁

Notice that here the cost share of agent 1 is negative, due to the convexity of c ! This may be considered as an undesirable feature of the cost sharing rule. Not only are costs increasing in the level of demand, in case of a convex cost function

each agent contributes to the negative externality. It seems fairly reasonable to demand a nonnegative contribution in those cases, so that none profits for just being there. The mainstream cost sharing literature includes positivity of cost shares into the very definition of a cost sharing rule. Here we will add it as a specific property:

Positivity μ is *positive* if $\mu(x, c) \geq 0_N$ for all (x, c) in its domain. All earlier discussed cost sharing rules have this property, except for the decreasing serial rule.

The decreasing serial rule is far more intuitive in case of economies of scale, in the presence of a concave cost function. The larger agents now are credited with a lower price per unit of the output good. Hougaard and Thorlund-Petersen (2001) and Koster (2002) propose variations on the serial rule that coincide with the increasing (decreasing) serial rule in case of a convex (concave) cost function, meeting the positivity requirement.

Marginal Pricing Rule

A popular way of pricing an output of a production facility is *marginal cost pricing*. The price of the output good is set to cover the cost producing one extra unit. It is frequently used in the domain of public services and utilities. However, a problem is that for concave cost functions, the method leads to budget deficits. An adapted form of marginal cost pricing splits these deficits equally over the agents. The *marginal pricing rule* is defined by

$$\begin{aligned}\mu_i^{\text{MP}}(x, c) &= x_i c'(x(N)) + \frac{1}{n} \\ &\quad \times [c(x(N)) - x(N)c'(x(N))].\end{aligned} \quad (20)$$

Note that in case of convex cost functions, agents can receive negative cost shares, just like it is the case with decreasing serial cost sharing.

Additive Cost Sharing and Rationing

The above cost sharing rules for homogeneous production models share the following properties:

Additivity $\mu(x, c_1 + c_2) = \mu(x, c_1) + \mu(x, c_2)$ for all relevant cost sharing problems. This property carries the same flavor as the homonymous property for cost games.

Constant Returns $\mu(x, c) = \vartheta x$ for linear cost functions c such that $c(y) = \vartheta y$ for all y . So if the agents do not cause any externality, the fixed marginal cost is taken as a price for the good.

It turns out that the class of all positive cost sharing rules with these properties can be characterized by solutions to *rationing problems* – which are the most basic of all models of distributive justice. A rationing problem among the agents in N consists of a pair $(x, t) \in \mathbb{R}_+^N \times \mathbb{R}_+$ such that $x(N) \geq t$; t is the available amount of some (in)divisible good and x is the set of demands. The inequality sees to the interpretation of rationing as not every agent may get all she wants.

A *rationing method* r is a solution to rationing problems, such that each problem (x, t) is assigned a vector of shares $r(x, t) \in \mathbb{R}_+^N$ such that $\mathbf{0}_N \leq r(x, t) \leq x$. The latter restriction is a weak solidarity statement assuring that everybody's demand be rationed in case of shortages.

For $t \in \mathbb{R}_+$ define the special cost function Γ_t by $\Gamma_t(y) = \min\{y, t\}$. The cone generated by these base functions lays dense in the space of all *absolutely continuous* cost functions c ; if we know what the values $\mu(x, \Gamma_t)$ are, then basically we know $\mu(x, c)$. Denote by $\mathcal{M}$ the class of all cost sharing rules with the properties positivity, additivity, and constant returns.

Theorem 7.3 *Moulin and Shenker* (Moulin 2002; Moulin and Shenker 1994)

Consider the following mappings associating rationing methods with cost sharing rules and vice versa:

$$r \mapsto \mu : \mu(x, c) = \int_0^{x(N)} c'(t) dr(x, t), \mu \mapsto r : r(x, t) = \mu(x, \Gamma_t).$$

These define an isomorphism between $\mathcal{M}$ and the space of all monotonic rationing methods.

So each monotonic rationing method relates to a cost sharing rule and vice versa. In this way μ^P is directly linked with the proportional rationing method, μ^{SR} to the *uniform gains* method, and μ^{SS} to the *random priority* method. Properties of rationing methods lead to properties of cost sharing rules and vice versa (Koster 2012).

Incentives in Cooperative Production

Stable Allocations, Stand-Alone Core

Suppose again, like in the framework of cooperative cost games, that (coalitions of) agents can decide to leave the cost sharing and organize their own production facility. Under the ability to replicate the technology, the question arises whether cost sharing rules induce stable cost share vectors.

Theorem 7.4 *Moulin* (1996)

For concave cost functions c , if μ is an anonymous and cross-monotonic cost sharing rule, then $\mu(x, c) \in \text{core}(c_x)$.

Under increasing returns to scale, this implies that μ^P, μ^{SR} are core selectors, but μ^{MC} is not. Tijs and Koster (1998) associates to each cost sharing problem (x, c) a *pessimistic* one, (x, c^*) ; here $c^*(y)$ reflects the maximum of marginal cost on $[0, x(N)]$ to produce y units,

$$c^*(y) = \begin{cases} \sup \left\{ \int_T c'(t) dt \mid T \subset [0, x(N)], \lambda(T) = y \right\} & \text{if } y \leq x(N), \\ c(y) & \text{else} \end{cases} \quad (21)$$

Here λ denotes the Lebesgue measure.

Theorem 7.5 *Koster* (2007)

For any cost sharing problem (x, c) , it holds $\text{core}(c_x^*) = \{\mu(x, c) \mid \mu \in \mathcal{M}\}$.

In particular this means that for $\mu \in \mathcal{M}$, it holds that $\mu(x, c) \in \text{core}(c_x)$ whenever c is concave, since this implies $c^* = c$. This result appeared earlier as a corollary to Theorem 7.3 (see Moulin 2002). Hougaard and Thorlund-Petersen (2001) and Koster (2002, 2012) show *nonlinear* cost

sharing rules yielding core elements for concave cost functions as well, so additivity is only a sufficient condition in the above statement. For average cost sharing, one can show more, $\mu^P(x, c) \in \text{core}(c_x)$ for all x precisely when the average cost $c(y)/y$ is decreasing in y .

Strategic Manipulation Through Reallocation of Demands

In the cooperative production model, there are other ways that agents may use to manipulate the final allocation. In particular, note that the serial procedure gives the larger demanders an advantage in case of positive externalities; as marginal costs decrease, the price paid by the larger agents per unit of output is lower than that of the smaller agents. In the other direction, larger demanders are punished if costs are convex. Then as the examples below show, this is just why the serial ideas are vulnerable to misrepresentation of demands, since combining demands and redistributing the output afterward can be beneficial.

Example 7.6 Consider the cost function c given by $c(y) = \min\{5y, 60 + 2y\}$. Such cost functions are part of daily life, whenever one has to decide upon telephone or energy supply contracts: usually customers get to choose between a contract with high fixed cost and a low variable cost and another with low or no fixed cost and a high variable price. Now consider the two cost sharing problems (x, c) and (x', c) , where $x = (10, 20, 30)$ and $x' = (0, 30, 30)$. The cost sharing problem (x', c) arises from (x, c) if agent 2 places a demand on behalf of agent 1 – without letting agent 3 know. The corresponding average and serial cost shares are given by

$$\begin{array}{ll} \mu^P(x, c) = (30, 60, 90) & \mu^P(x', c) = (0, 90, 90) \\ \mu^{\text{SR}}(x, c) = (40, 60, 80) & \mu^{\text{SR}}(x', c) = (0, 90, 90) \end{array}$$

Notice that the total of average cost shares for 1 and 2 is the same in both cost sharing problems. But if the serial rule were used, these agents can profit by *merging* their demands; if agent 2 returns agent 1's demand and requires a payment from

agent 1 between €30 and €40, then both agents will have profited by such *merging* of demand.

Example 7.7 Consider the 5-agent cost sharing problems (x, c) and $(\bar{x}, c)$ with $x = (1, 2, 3, 0, 0)$, $\bar{x} = (1, 2, 1, 1, 1)$ and convex cost function $c(y) = \frac{1}{2}y^2$. $(\bar{x}, c)$ arises out of (x, c) if agent 3 splits her demand over agents 4 and 5 as well. Then

$$\begin{array}{l} \mu^P(x, c) = (6, 12, 18, 0, 0) \\ \mu^P(\bar{x}, c) = (6, 12, 6, 6, 6), \\ \mu^{\text{SR}}(x, c) = (3, 11, 22, 0, 0) \\ \mu^{\text{SR}}(\bar{x}, c) = (5, 16, 5, 5, 5). \end{array}$$

The aggregate of average cost shares for agents 3, 4, and 5 does not change. But notice that according to the serial cost shares, there is a clear advantage for the agents. Instead of paying 22 in the original case, now the total of payments equals 15. Agent 3 may consider a transfer between 0 and 7 to 3 and 4 for their collaboration and still be better off. In general, in case of a convex cost function, the serial rule is vulnerable with respect to manipulation of demands through splitting.

◁

Note that in the above cases, the proportional cost sharing rule does prescribe the same cost shares. It is a non-manipulable rule: reshuffling of demands will not lead to different aggregate cost shares. The rule does not discriminate between units, when a unit produced is irrelevant. It is actually a very special feature of the cost sharing rule that is basically not satisfied by any other cost sharing rule.

Theorem 7.8 Assume that N contains at least three agents. The proportional cost sharing rule is the unique rule that charges nothing for a null demand and meets any one of the following properties:

- Independence of merging and splitting
- No advantageous reallocation
- Irrelevance of reallocation

The second property shows even a stronger property than merging and splitting in that agents may redistribute the demands in any preferred way without changing the aggregate cost shares of the agents involved. The third property says that in such cases, the cost shares of the other agents do not change. Then this makes proportional cost sharing compelling in situations where one is not capable of detecting the true demand characteristics of individuals.

Demand Games for $\mathcal{P}^1$

Consider demand games $G(\mu, c)$ as in (11), section “Demand Games,” where now μ is a cost sharing rule on $\mathcal{P}^1$. These games with uncountable strategy spaces are more complex than the demand games that we studied before.

The set of consequences for players is now given by $C = \mathbb{R}_+^2$, combinations of levels of production and costs (see section “Strategic Demand Games”). Then an individual i ’s preference relation is *convex* if for the corresponding utility function u_i and all pairs $z, z' \in C$ it holds

$$\begin{aligned} u_i(z) = u_i(z') &\Rightarrow u(tz + (1-t)z) \\ &\geq u_i(z) \text{ for all } t \in [0, 1]. \end{aligned} \quad (22)$$


This means that a weighted average of the consequences is weakly preferred to both

consequences, if these are equivalent. Such utility functions u_i are called *quasi-concave*. An example of convex preferences are those related to *linear utility* functions of type $u_i(x, y) = \alpha x - y$. Moreover, *strictly convex* preferences are those with strict inequality in (22) for $0 < t < 1$; the corresponding utility functions are *strictly quasi-concave*. Special classes of preferences are the following:

- $\mathcal{L}$: the class of all convex and continuous preferences utility functions that are non-decreasing in the service component x , non-increasing in the cost component y , non-locally satiated, and decreasing on $(x, c(x))$ for x large enough. The latter restriction is no more than assuring that agents will not place requests for unlimited amounts of the good (Fig. 8).
- $\mathcal{L}^*$: the class of *bi-normal* preferences in $\mathcal{L}$. Basically, if such a preference is represented by a differentiable utility function u , then the slope dy/dx of the indifference contours is non-increasing in x and non-decreasing in y . For a concise definition, see Watts (2002). Examples include Cobb-Douglas utility functions and also those of type $u_i(x, y) = \alpha(x) - \beta(y)$ where α and β are concave and convex functions, respectively. A typical plot of level curves of such utility functions is in Fig. 9. Note that the


Cost Sharing in Production Economies,

Fig. 8 Linear, convex preferences, $u(x, y) = 2x - y$. The contours indicate indifference curves, i.e., sets of type $\{(x, y) | u(x, y) = k\}$, the k -level curve of u



Cost Sharing in Production Economies,

Fig. 9 Strictly convex preferences, $u(x, y) = \sqrt{x} - e^{0.5y}$. The straight line connecting any two points on the same contour lies in the lighter area – with higher utility. Fix the y -value, then an increase of x yields higher utility, whereas for fixed x an increase of y causes the utility to decrease



approach differs from the standard literature where agents have preferences over endowments. Here costs are “negative” endowments. In the latter interpretation, the condition can be read as that the *marginal rate of substitution* is nonpositive. At equal utility, an increase of the level of output has to be compensated by a decrease in the level of input good.

Nash Equilibria of Demand Games in a Simple Case

Consider a production facility shared by three agents $N = \{1, 2, 3\}$ with cost function $c(y) = \frac{1}{2}y^2$. Assume that the agents have quasi-linear utilities in $\mathcal{L}$, i.e., $u_i(x_i, y_i) = \alpha_i x_i - y_i$ for all pairs $(x_i, y_i) \in \mathbb{R}_+^2$. Below the Nash equilibrium in the serial and proportional demand game is calculated in two special cases. This numerical example is based on Moulin and Shenker (1992).

Proportional Demand Game

Consider the corresponding proportional demand game, $G(\mu^P, c)$, with utility over actions given by

$$\begin{aligned} U_i^P(x) &= \alpha_i x_i - \mu^P(x, c) \\ &= \alpha_i x_i - \frac{1}{2} x_i x(N). \end{aligned} \quad (23)$$

In a Nash equilibrium x^* of $G(\mu^P, c)$, each player i gives a best response on x_{-i}^* , the action profile of the other agents. That is, player i chooses $x_i^* \in \arg \max_i U_i^P(t, x_{-i}^*)$. Then first-order conditions imply for an interior solution

$$\alpha_i - \frac{1}{2} x^*(N) - \frac{1}{2} x_i^* = 0 \quad (24)$$

for all $i \in N$. Then $x^*(N) = \frac{1}{2}(\alpha_1 + \alpha_2 + \alpha_3)$ and $x_i^* = 2\alpha_i - \frac{1}{2}(\alpha_1 + \alpha_2 + \alpha_3)$.

Serial Demand Game

Consider the same production facility and the demand game $G(\mu^{SR}, c)$, corresponding to the serial rule. Then the utilities over actions are given by

$$U_i^{SR}(x) = \alpha_i x_i - \mu^{SR}(x, c). \quad (25)$$

Now suppose $\bar{x}$ is a Nash equilibrium of this game, and assume without loss of generality that $\bar{x}_1 \leq \bar{x}_2 \leq \bar{x}_3$. Then player 1 with the smallest equilibrium demand maximizes the expression

$$U_1^{SR}((t, \bar{x}_2, \bar{x}_3)) = \alpha_1 t - c(3t)/3 = \alpha_1 t - \frac{3}{2} t^2$$

at $\bar{x}_1$, from which we may conclude that $\alpha_1 = 3\bar{x}_1$. In addition, in equilibrium, player 2, maximizes

$$U_2^{\text{SR}}(\bar{x}_1, t, \bar{x}_3) = \alpha_2 t - \left(\frac{1}{3} c(3\bar{x}_1) + \frac{1}{2} (c(\bar{x}_1 + 2t) - c(3\bar{x}_1)) \right),$$

for $t \geq \bar{x}_1$, yielding $\alpha_2 = \bar{x}_1 + 2\bar{x}_2$. Finally, the equilibrium condition for player 3 implies $\alpha_3 = \bar{x}(N)$. Then it is not hard to see that actually this constitutes *the* serial equilibrium.

Comparison of Proportional and Serial Equilibria (I)

Now let's compare the serial and the proportional equilibrium in the following two cases:

- (i) $\alpha_1 = \alpha_2 = \alpha_3 = \alpha$
- (ii) $\alpha_1 = \alpha_2 = 2, \alpha_3 = 4$.

Case (i): We get $x^* = (\frac{1}{2}\alpha, \frac{1}{2}\alpha, \frac{1}{2}\alpha)$ and $\bar{x}_i = (\frac{1}{3}\alpha, \frac{1}{3}\alpha, \frac{1}{3}\alpha)$ for all i . The resulting equilibrium payoff vectors are given by

$$U^{\text{P}}(x^*) = \left(\frac{1}{8}\alpha^2, \frac{1}{8}\alpha^2, \frac{1}{8}\alpha^2 \right) \text{ and } U^{\text{SR}}(\bar{x}) = \left(\frac{1}{6}\alpha^2, \frac{1}{6}\alpha^2, \frac{1}{6}\alpha^2 \right).$$

Not only the average outcome is less efficient than its serial counterpart, it is also Pareto-inferior to the latter.

Case (ii): The proportional equilibrium is a boundary solution, $x^* = (0, 0, 2)$, with utility profile $U^{\text{P}}(x^*) = (0, 0, 8)$. The serial equilibrium strategies and utilities are given by

$$\bar{x} = \left(\frac{2}{3}, \frac{2}{3}, \frac{8}{3} \right), U^{\text{SR}}(\bar{x}) = \left(\frac{2}{3}, \frac{2}{3}, 4 \right).$$

Notice that the serial equilibrium is now less efficient, but is not Pareto dominated by the proportional utility distribution.

Uniqueness of Nash Equilibria in P¹-Demand Games

In the above demand games, there is a unique Nash equilibrium which serves as a prediction of actual play. This needs not hold for any game. In the literature strategic characterizations of cost sharing rules are discussed in terms of uniqueness

of equilibrium in the induced cost games, relative to specific domains of preferences and cost functions. Below we will discuss the major findings of Watts (2002). These results concern a broader cost sharing model with the notion of a cost function as a differentiable function $\mathbb{R}_+ \rightarrow \mathbb{R}_+$. So in this paragraph, such cost function can decrease; fixed cost need not be 0. This change in setup is not crucial to the overall exposition since the characterizations below are easily interpreted within the context of $\mathcal{P}^1$.

Demand Monotonicity The mapping $t \mapsto \mu_i(t, x_{-i}, c)$ is non-decreasing; μ is strictly demand monotonic if this mapping is increasing whenever c is increasing.

Smoothness The mapping $x \mapsto \mu(x, c)$ is continuously differentiable for all continuously differentiable $c \in \mathcal{C}^1$.

Recall that $\mathcal{L}^*$ is the domain of all binormal preferences.

Theorem 7.9 Watts (2002)

Fix a differentiable cost function c and a demand monotonic and smooth cost sharing rule μ . A cost sharing game $G(\mu, c)$ has a unique equilibrium whenever agents' preferences belong to $\mathcal{L}^$, only if, for all $x = (x_1, \dots, x_n)$*

- *Every principal minor of the matrix W with rows w^i is nonnegative for all.*

$$w^i \in \left\{ \left(\frac{\partial \mu_i}{\partial x_1}, \dots, \frac{\partial \mu_i}{\partial x_n} \right), \left(\frac{\partial^2 \mu_i}{\partial x_1 \partial x_1}, \dots, \frac{\partial^2 \mu_i}{\partial x_i \partial x_n} \right) \right\} \quad (26)$$

- *The determinant of the Hessian matrix corresponding to the mapping $x \mapsto \mu(x, c)$ is strictly positive.*

Sufficient condition to have uniqueness of equilibrium is that the principle minor of the matrix W is strictly positive.

The impact of this theorem is that one can characterize the class of cost functions yielding unique equilibria if the domain of preferences is $\mathcal{L}^*$.

- $G(\mu^{\text{sr}}, c)$, $G(\mu^{\text{dsr}}, c)$: Necessary condition for uniqueness of equilibrium is that c is strictly convex, i.e., $c'' > 0$. Sufficient condition is that c is increasing and strictly convex. Actually, Watts (2002) also shows that the conclusions for the serial rule do not change when $\mathcal{L}$ is used instead of $\mathcal{L}^*$. As will get more clearer below, the serial games have unique strategic properties.
- $G(\mu^{\text{p}}, c)$: The necessary and sufficient conditions are those for the serial demand game, including $c'(y) > c(y)/y$ for all $y \neq 0$. Notice that the latter property does not pose additional restrictions on cost functions within the framework of P^1 .
- $G(\mu^{\text{ss}}, c)$: Necessary condition is $c'' > 0$. In general it is hard to establish uniqueness if more than two players are involved.
- $G(\mu^{\text{mp}}, c)$: Even in two-player games, uniqueness is not guaranteed. For instance, uniqueness is guaranteed for cost functions $c(y) = y^\alpha$ only if $1 < \alpha \leq 3$. For $c(y) = y^4$, there are preference profiles in $\mathcal{L}^*$ such that multiple equilibria reside.

Decreasing Returns to Scale

The above theorem basically shows that uniqueness of equilibrium in demand games related to $\mathcal{P}^1$ can be achieved for preferences in $\mathcal{L}^*$ if only costs are convex, i.e., the technology exhibits decreasing returns to scale. Starting point in the literature to characterize cost sharing rules in terms of their strategic properties is the seminal paper by Moulin and Shenker (1992). Their finding is that on $\mathcal{L}$ basically μ^{sr} is the only cost sharing rule of which the corresponding demand game passes the unique equilibrium test like in Theorem 7.9.

Call a smooth and strictly demand monotonic cost sharing rule μ *regular* if it is *anonymous*, so that the name of an agent does not have an impact on her cost share.

Theorem 7.10 Moulin and Shenker (1992)

Let c be a strictly convex continuously differentiable cost function, and let μ be a regular cost sharing rule. The following statements are equivalent:

- $\mu = \mu^{\text{sr}}$.
- For all profiles $(u_1, u_2, \dots, u_n)$ of utilities in $\mathcal{L}$, $G(\mu, c)$ has at most one Nash equilibrium.
- For all profiles $(u_1, u_2, \dots, u_n)$ of utilities in $\mathcal{L}$, every Nash equilibrium of $G(\mu, c)$ is also a strong equilibrium, i.e., no coalition can coordinate in order to improve the payoff for all its members.

This theorem makes a strong case for the serial cost sharing rule, especially when one realizes that the serial equilibrium is the unique element surviving successive elimination of strictly dominated strategies. Then this equilibrium may naturally arise through evolutive or educative behavior; it is a robust prediction of noncooperative behavior. Recent experimental studies are in line with this theoretical support (see Chen 2003; Razzolini et al. 2004). Proposition 1 in Watts (2002) shows how easy it is to construct preferences in $\mathcal{L}$ such that regular cost sharing rules other than μ^{sr} give rise to multiple equilibria in the corresponding demand game, even in two-agent cost sharing games.

Besides other fairness concepts in the distributive literature, the most compelling is *envy-freeness*. An allocation passes the no-envy test if no player prefers her own allocation less than that of other players.

No Envy Test Let x be a demand profile and y a vector of cost shares. Then the allocation $(x_i, y_i)_{i \in N}$ is *envy-free* if for all $i, j \in N$ it holds $u_i(x_i, y_i) \geq u_i(x_j, y_j)$.

It is easily seen that the allocations associated with the serial equilibria are all envy-free.

Increasing Returns to Scale

As Theorem 7.9 already shows, uniqueness of equilibrium in demand games for all utility profiles in $\mathcal{L}^*$ is in general inconsistent with concave cost functions.

Theorem 7.11 de Frutos (1998)

Let c be a strictly concave continuously differentiable cost function, and let μ be a regular cost sharing rule. The following statements are equivalent:

- $\mu = \mu^{\text{DSR}}$ or $\mu = \mu^{\text{SR}}$.
- For all utility profiles $u = (u_i)_{i \in N}$ in $\mathcal{L}$, the induced demand game $G(\mu, c)$ has at most one Nash equilibrium.
- For all utility profiles $u = (u_i)_{i \in N}$ in $\mathcal{L}$, every Nash equilibrium of the game $G(\mu, c)$ is a strong Nash equilibrium as well.

Moreover, if the curvature of the indifference curves is bigger than that of the curve generated by the cost sharing rule as in Fig.10, then the second and third statements are equivalent to $\mu = \mu^{\text{DSR}}$.

Theorem 7.12 *Moulin (1996)*

Assume agents have preferences in $\mathcal{L}^*$. The serial cost sharing rule is the unique continuous, cross-monotonic, and anonymous cost sharing rule for which the Nash equilibria of the corresponding demand games all pass the no-envy test.

Comparison of Serial and Proportional Equilibria (II)

Just as in the earlier analysis in section “[Efficiency and Strategy-Proof Cost Sharing Mechanisms](#),” performance of cost sharing rules can be

measured by the related surpluses in the Nash equilibria of the corresponding demand games. Assume in this section that the preferences of the agents are quasi-linear in cost shares and represented by functions $U_i(x_i, y_i) = u_i(x_i) - y_i$. Moreover, assume that u_i is non-decreasing and concave, $u_i(0) = 0$. Then the surplus at the demand profile x and utility profile is the number $\sum_{i \in N} u_i(x_i) - c(x(N))$. Define the *efficient surplus* or *value* of N relative to c and U by

$$v(c, U) = \sup_{x \in \mathbb{R}_+^N} \sum_{i \in N} u_i(x_i) - c(x(N)). \quad (27)$$


Denote the set of Nash equilibria in the demand game $G(\mu, c)$ with profile of preferences U by $\text{NE}(\mu, c, U)$. Given c, μ the *guaranteed (relative) surplus* of the cost sharing rule μ for N is defined by

$$\gamma(c, \mu) = \inf_{U, x \in \text{NE}(\mu, c, U)} \frac{\sum_{i \in N} u_i(x_i) - c(x(N))}{v(c, U)}. \quad (28)$$

Here the infimum is taken over all utility profiles discussed above. This measure is also called the *price of anarchy* of the game (see Koutsoupas and

Cost Sharing in Production Economies,

Fig. 10 Scenario (ii). The indifference curves of agent 1 together with the curve $\kappa : t \mapsto \mu_1^{\text{SR}}((t, \bar{x} - 1), c)$. Best response of player 1 against $\bar{x}_{-1} = (\frac{2}{3}, \frac{8}{3})$ is the value of x where the graph of κ is tangent to an indifference curve of u_1



Papadimitriou 1999). Let C^* be the set of all convex increasing cost functions with $\lim_{y \rightarrow \infty} c(y)/y = \infty$. Then Moulin (2010) shows that for the serial and the proportional rule, the guaranteed surplus is at least $1/n$. But sometimes the distinction is eminent. Define the number

$$\delta(y) = \frac{yc''(y)}{c'(y) - c'(0)},$$

which is a kind of elasticity. The below theorem shows that on certain domains of cost functions with bounded δ , the serial rule prevails over the proportional rule. For large n the guaranteed surplus at μ^{SR} is of order $1/\ln(n)$, that of μ^{AV} of order $1/n$. Write $K_n = 1 + \frac{1}{3} + \dots + \frac{1}{2n-1} \approx 1 + \frac{\ln n}{2}$, then we obtain the following.

Theorem 7.13 Moulin (2010)

For any convex increasing cost function c with $\lim_{y \rightarrow \infty} c(y)/y = \infty$, it holds that

- If c' is concave and $\inf\{\delta(y) \mid y \geq 0\} = p > 0$, then

$$\gamma(c, \mu^{\text{SR}}) \geq \frac{p}{K_n} \quad \gamma(c, \mu^{\text{AV}}) \leq \frac{4}{n+3}$$

- If c' is convex and $\sup\{\delta(y) \mid y \geq 0\} = p < \infty$, then

$$\begin{aligned} \gamma(c, \mu^{\text{SR}}) &\geq \frac{1}{2p-1} \frac{1}{K_n} \quad \frac{4}{n+3} \leq \gamma(c, \mu^{\text{AV}}) \\ &\leq \frac{4(2p-1)}{n} \end{aligned}$$

A Word on Strategy-Proofness in P^1

Recall the discussion on strategy-proofness in section “[Strategy-Proofness](#).” The serial demand game has a unique strong Nash equilibrium in case costs are convex and preferences are drawn from $\mathcal{L}$. Suppose the social planner aims at designing a mechanism to implement the outcomes associated with these equilibria. Moulin and Shenker (1992) shows an efficient way to

implement this *serial choice function* by an *indirect* mechanism. It is defined through a multi-stage game which mimics the way the serial Nash equilibria are calculated. It is easily seen that here each agent has a unique dominant strategy, in which demands result from optimization of the true preferences. Then this gives rise to a strategy-proof mechanism.

Note that the same approach cannot be used for the proportional rule. The strategic properties of the proportional demand game are weaker than that of the serial demand game in several aspects. First of all, it is not hard to find preference profiles in $\mathcal{L}$ leading to multiple equilibria. Whereas uniqueness of equilibrium can be repaired by restricting $\mathcal{L}$ to $\mathcal{L}^*$, the resulting equilibria are in general not strong (like the serial counterparts). In the proportional equilibria, there is over-production; see, e.g., the example in section “[Proportional Demand Game](#)” where a small uniform reduction of demands yields higher utility for all the players. Besides, a single-valued Nash equilibrium selection corresponds to a strategy-proof mechanism *provided* the underlying domain of preferences is rich, and $\mathcal{L}^*$ is not. Though richness is not a necessary condition, the proportional rule is not consistent with a strategy-proof demand game.

Bayesian P^1 -Demand Games

Recall that at the basis of a strategic game, there is the assumption that each player *knows* all the ingredients of the game. However, as Kolpin and Wilbur (2005) argues, production cost and output quality may vary unpredictably as a consequence of the technology and worker quality. Besides that, changes in the available resources and demands will have not foreseen influences on individual preferences. On top of that, the players may have asymmetrical information regarding the nature of uncertainty. Kolpin and Wilbur (2005) study the continuous homogeneous cost sharing problem within the context of a *Bayesian demand game* (Harsanyi 1967), where these uncertainties are taking into account. The qualifications of the serial rule in the stochastic model are roughly the same as in the deterministic framework.

Continuous Heterogeneous Output Model, $\mathcal{P}^n$

The analysis of continuous cost sharing problems for multiservice facilities is far more complex than the single-output model. The literature discusses two different models, one where each agent i demands a different good and one where agents may require mixed bundles of goods. As the reader will notice, the modeling and analysis of solutions differs in abstraction and complexity. In order to concentrate on the main ideas, here we will stick to the first model, where goods are identified with agents. This means that a demand profile is a vector $x \in \mathbb{R}_+^N$, where x_i denotes the demand of agent i for good i . From now we deal with technologies described by continuously differentiable cost functions $c : \mathbb{R}_+^N \rightarrow \mathbb{R}_+$, non-decreasing, and $c(0_N) = 0$; the class of all such functions is denoted $\mathcal{C}^n$.

Extensions of Cost Sharing Rules

The single-output model is connected to the multi-output model via the homogeneous cost sharing problems. Suppose that for $c \in \mathcal{C}^n$ there is a function $c_0 \in \mathcal{C}$ such that $c(x) = c_0(x(N))$ for all x . For instance, such functions are found if we distinguish between production of blue and red cars; the color of the car does not affect the total production costs. Essentially, a homogeneous cost sharing problem (x, c) may be solved as if it were in $\mathcal{P}^1$. If μ is the compelling solution on $\mathcal{P}^1$, then any cost sharing rule on $\mathcal{P}^n$ should determine the same solution on the class of *homogeneous* problems therein. Formally the cost sharing rule on $\mathcal{P}^n$ extends $\bar{\mu}$ on $\mathcal{P}^1$ if for all homogeneous cost sharing problems (x, c) it holds $\bar{\mu}(x, c) = \mu(x, c_0)$. In general a cost sharing rule μ on $\mathcal{P}^1$ allows for a whole class of extensions. Below we will focus on extensions of μ^{SR} , μ^{P} , μ^{SS} .

Measurement of Scale

Around the world quantities of goods are measured by several standards. Length is expressed in inches or centimeters, volume in gallons or liters, and weight in ounces to kilos. Here measurement conversion involves no more than multiplication with a fixed scalar. When such linear scale conversions do not have any effect on final cost shares, a cost sharing rule is called *scale invariant*.

It is an *ordinal* rule if this invariance extends to essentially all transformations of scale. Scale invariance captures the important idea that the relative cost shares should not change, whether we are dividing 1 euro or 1000 euros. Ordinality may be desirable, but for many purposes too strong as a basic requirement. Formally, a *transformation of scale* is a mapping $f : \mathbb{R}_+^N \rightarrow \mathbb{R}_+^N$ such that $f(x) = (f_1(x_1), f_2(x_2), \dots, f_n(x_n))$ for all x and each of the coordinate mappings f_j is differentiable and strictly increasing.

Ordinality

A cost sharing rule μ on $\mathcal{P}^n$ is *ordinal* if for all transformations of scale f and all cost sharing problems $(x, c) \in \mathcal{P}^n$ it holds that

$$\mu(x, c) = \mu(f(x), c \circ f^{-1}). \quad (29)$$

Scale Invariance

A cost sharing rule μ on $\mathcal{P}^n$ is *scale invariant* if (29) holds for all linear transforms f . Under a scale invariant cost sharing rule, final cost shares do not change by changing the units in which the goods are measured.

Path-Generated Cost Sharing Rules

Many cost sharing rules on $\mathcal{P}^n$ calculate the cost shares for $(x, c) \in \mathcal{P}^n$ by the total of marginal costs along some *production path* from 0 toward x . Here a path for x is a non-decreasing mapping $\gamma^x : \mathbb{R}_+ \rightarrow \mathbb{R}_+^N$ such that $\gamma(0) = 0_N$ and there is a $T \in \mathbb{R}_+$ with $\gamma^x(T) = x$. The cost sharing rule *generated by the path collection* $\gamma = \{\gamma^x | x \in \mathbb{R}_+^N\}$ is defined by

$$\mu_i^\gamma(x, c) = \int_0^\infty \partial_i c(\gamma^x(t)) (\gamma_i^x)'(t) dt. \quad (30)$$

Special path-generated cost sharing rules are the *fixed-path cost sharing rules*; a single path $\gamma^* : \mathbb{R}_+ \rightarrow \mathbb{R}_+^N$ with the property that $\lim_{t \rightarrow \infty} \gamma_i^*(t) = \infty$ defines the whole underlying family of paths. More precisely, the *fixed-path cost sharing rule* μ generated by γ^* is the path-generated rule for the family of paths $\{\gamma^x | x \in \mathbb{R}_+^N\}$ defined by $\gamma^x(t) = \gamma^*(t) \wedge x$, the vector with coordinates $\min\{\gamma_i^*(t), x_i\}$. So

the paths are no more than the projections of $\gamma^*(t)$ on the cube $[0, x]$. Below we will see many examples of (combinations of) such fixed-path methods.

Aumann-Shapley Rule

The early characterizations by Billera and Heath (1982) and Mirman and Tauman (1982) on this rule set off a vast growing literature on cost sharing models with variable demands. Billera et al. (1978) suggested to use the Aumann-Shapley rule to determine telephone billing rates in the context of sharing the cost of a telephone system. This extension of proportional cost sharing calculates marginal costs along the path $\gamma^{\text{AS}}(t) = tx$ for $t \in [0, 1]$. Then

$$\mu_i^{\text{AS}}(x, c) = x_i \int_0^1 \partial_i c(tx) dt. \quad (31)$$

The Aumann-Shapley rule can be interpreted as the Shapley value of the *non-atomic* game where each unit of the good is a player (see Aumann and Shapley 1974). It is the uniform average over marginal costs along all increasing paths from 0_N to x . The following is a classic result in the cost sharing literature:

Theorem 7.14 *Mirman and Tauman (1982), Billera and Heath (1982)*

There is only one additive, positive, and scale invariant cost sharing rule on $\mathcal{P}^n$ that extends the proportional rule, and this is μ^{AS} .

Example 7.15 If c is positively homogeneous, i.e., $c(\alpha y) = \alpha c(y)$ for $\alpha \geq 0$ and all $y \in \mathbb{R}_+^N$, then

$$\mu_i^{\text{AS}}(x, c) = \frac{\partial_i c}{\partial x_i}(x),$$

i.e., μ^{AS} calculates the marginal costs of the i -th good at the final production level x . The risk measures (cost functions) as in Denault (2001) are of this kind. $\triangleleft$

Friedman-Moulin Rule

This serial extension (Friedman and Moulin 1999) calculates marginal costs along the diagonal path, i.e., $\gamma^{\text{FM}}(t) = t1_N \wedge x$

$$\mu_i^{\text{FM}}(x, c) = \int_0^{x_i} \partial_i c(\gamma^x(t)) dt \quad (32)$$

This fixed-path cost sharing rule is demand monotonic. As far as invariance with the choice of unit is concerned, the performance is bad as it is not a scale invariant cost sharing rule.

Moulin-Shenker Rule

This fixed-path cost sharing rule is proposed as an ordinal serial extension by Sprumont (1998). Suppose that the partial derivatives of $c \in \mathcal{C}^n$ are bounded away from 0, i.e., there is a such that $\partial_i c(x) > a$ for all $x \in \mathbb{R}_+^N$. The Moulin-Shenker rule μ^{MS} is generated by the path γ^{MS} as solution to the system of ordinary differential equations

$$\gamma'_i(t) = \begin{cases} \frac{\sum_{j \in N} \partial_j c(\gamma(t))}{\partial_i c(\gamma(t))} & \text{if } \gamma_i(t) < x_i, \\ 0 & \text{else.} \end{cases} \quad (33)$$

The interpretation of this path is that at each moment the total expenditure for production of extra units for the different agents is equal; if good 2 is twice as expensive as good 1, then the production device γ^{MS} will produce twice as much of the good 1. The serial rule embraces the same idea – as long as an agent desires extra production, the corresponding incremental costs are equally split. This makes μ^{MS} a natural extension of the serial rule. Call t_i^* the completion time of production for agent i , i.e., $\gamma_i^{\text{MS}}(t) < x_i$ if $t < t_i^*$ and $\gamma_i^{\text{MS}}(t_i^*) = x_i$. Assume without loss of generality that these completion times are ordered such that $0 = t_0^* \leq t_1^* \leq t_2^* \leq \dots \leq t_n^*$, then the Moulin-Shenker rule is given by

$$\mu_i^{\text{MS}}(x, c) = \sum_{\ell=1}^i \frac{c(\gamma^{\text{MS}}(t_\ell^*)) - c(\gamma^{\text{MS}}(t_{\ell-1}^*))}{n - \ell + 1}. \quad (34)$$

Note that the path varies with the cost function and that this is the reason why μ^{MS} is a non-additive solution. Such solutions – though intuitive – are in general notoriously hard to analyze. There are two

axiomatic characterizations of the Moulin-Shenker rule. The first is by Sprumont (1998), in terms of the *serial principle* and the technical condition that a cost sharing rule be a partial differentiable functions of the demands. The other characterization by Koster (2007) is more in line with the ordinal character of this *serial extension*.

Continuity

A cost sharing rule μ on P^n is *continuous* if $q \mapsto \mu(q, c)$ is continuous on $\mathbb{R}_+^N$ for all c .

Continuity is weaker than partial differentiability, as it requires stability of the solution with respect to small changes in the demands.

Upper Bound

A cost sharing rule μ satisfies *upper bound* if for all $(q, c) \in \mathcal{P}^n, i \in N$

$$\mu_i(q, c) \leq \max_{y \in [0, q]} \partial_i c(y).$$

Upper bound provides each agent with a conservative and rather pessimistic estimate of her cost share, based on the maximal value of the corresponding marginal cost toward the aggregate demand.

Suppose that d is a demand profile smaller than q . A reduced cost sharing problem is defined by $(q - d, c^d)$ where c^d is defined by $c^d(y) = c(y + d) - c(d)$. So c^d measures the incremental cost of production beyond the level d .

Self-Consistency

A cost sharing rule μ is *self-consistent* if for all cost sharing problems $(q, c) \in \mathcal{P}^n$ with $q_N \setminus S = 0_{N \setminus S}$ for some $S \subseteq N$, and $d \leq q$ such that $\mu_i(d, c) = \mu_j(d, c)$ for all $\{i, j\} \subset S$, then $\mu(q, c)_S = \mu(d, c)_S + \mu(q - d, c^d)_S$.

So, self-consistency expresses the idea that if cost shares of agents with nonzero demand differ, then this is not due to the part of the problem that they are equally charged for but due to the asymmetries in the related reduced problem. The

property is reminiscent of the *step-by-step negotiation property* in the bargaining literature (see Kalai 1977).

Theorem 7.16 Koster (2007)

There is only one continuous, self-consistent, and scale invariant cost sharing rule satisfying upper bounds, which is the Moulin-Shenker rule.

Shapley-Shubik Rule

For each demand profile x , the stand-alone cost game c_x is defined as before. Then the Shapley-Shubik rule is no more than the Shapley value of this game, i.e., $\mu^{\text{SS}}(x, c) = \Phi(c_x)$. The Shapley-Shubik rule is ordinal.

A Numerical Example

Consider the cost sharing problem (x, c) with $N = \{1, 2\}$, $x = (5, 10)$, and $c \in \mathcal{C}^2$ is given by $c(t_1, t_2) = e^{2t_1+t_2} - 1$ on $[0, 10] \times [0, 10]$. We calculate the partial derivatives $\partial_1 c(t_1, t_2) = 2e^{2t_1+t_2} = 2\partial_2 c(t_1, t_2)$ for all $(t_1, t_2) \in \mathbb{R}_+^2$. The Aumann-Shapley path is given by $\gamma(t) = (5t, 10t)$ for $t \in [0, 1]$ and

$$\begin{aligned} \mu_1^{\text{AS}}(x, c) &= \int_0^1 \partial_1 c(5t, 10t) \cdot 5 dt \\ &= \int_0^1 2e^{20t} \cdot 5 dt = \frac{1}{2} (e^{20} - 1) \\ \mu_2^{\text{AS}}(x, c) &= \int_0^1 \partial_2 c(5t, 10t) \cdot 10 dt \\ &= \int_0^1 e^{20t} \cdot 10 dt = \frac{1}{2} (e^{20} - 1) \end{aligned}$$

The Friedman-Moulin rule uses the path

$$\begin{aligned} \gamma^{\text{FM}}(t) &= t1_N \wedge q \\ &= \begin{cases} (t, t) & \text{if } 0 \leq t < 5, \\ (5, 5 + t) & \text{if } 5 \leq t, \end{cases} \end{aligned}$$

and the corresponding cost shares are calculated as follows:

$$\begin{aligned}\mu_1^{\text{FM}}(x, c) &= \int_0^5 \partial_1 c(t, t) dt \\ &= \int_0^5 2e^{3t} dt = \frac{2}{3}(e^{15} - 1), \\ \mu_2^{\text{FM}}(x, c) &= \int_0^5 \partial_2 c(t, t) dt + \int_5^{10} \partial_2 c(5, 5 + t) dt \\ &= \frac{1}{3}(e^{15} - 1) + e^{20} - e^{15}.\end{aligned}$$

Note that both discussed cost sharing rules use one and the same path for all cost sharing problems with demand profile x . This is characteristic for *additive* cost sharing rules (see, e.g., Friedman 2004; Haimanko 2000).

Now turn to the Moulin-Shenker rule. Since $\partial_1 c = 2\partial_2 c$ everywhere on $[0, 10] \times [0, 10]$, according to the solution γ^{MS} of (33), until one of the demands is reached, for each produced unit of good 1, two units of good 2 are produced. In particular there is a parametrization $\bar{\gamma}$ of γ^{MS} such that $\bar{\gamma}(t) = (t, 2t)$ for $0 \leq t \leq 5$. The corresponding cost shares are equal since $\bar{\gamma}$ reaches both coordinates of x at the same time, so $\mu_1^{\text{MS}}(x, c) = \mu_2^{\text{MS}}(x, c) = \frac{1}{2}c(\bar{\gamma}(5)) = \frac{1}{2}c(5, 10) = \frac{1}{2}(e^{20} - 1)$. Now suppose that the demands are summarized by $x^* = (10, 10)$. In order to calculate $\mu^{\text{MS}}(x^*, c)$, notice that there is a parametrization of γ^* of the corresponding path γ^{MS} such that (Figs. 11 and 12)

$$\gamma^*(t) = \begin{cases} (t, 2t) & \text{if } t \leq 5, \\ (t, 10) & \text{for } 5 < t \leq 10, \end{cases}$$

Notice that this path extends $\bar{\gamma}$ just to complete service for agent 1, so that – like before – agent 2 only contributes while $t < 5$. Then the cost shares are given by


$$\begin{aligned}\mu_2^{\text{MS}}(x^*, c) &= \frac{1}{2}c(\gamma^*(5)) = \frac{1}{2}c(5, 10) \\ &= \frac{1}{2}(e^{20} - 1), \\ \mu_1^{\text{MS}}(x^*, c) &= \mu_2^{\text{MS}}(x^*, c) + c(\gamma^*(10)) - c(\gamma^*(5)) \\ &= e^{30} - \frac{1}{2}(e^{20} + 1).\end{aligned}$$

For x^* the cost sharing rules μ^{AS} and μ^{FM} use essentially the same symmetric path $\gamma(t) = (t, t)$,


so that it is easily calculated that $\mu^{\text{AS}}(x^*, c) = \mu^{\text{FM}}(x^*, c) = (\frac{2}{3}(e^{30} - 1), \frac{1}{3}(e^{30} - 1))$.

Axiomatic Characterization of Fixed-Path Rules

Recall demand monotonicity as a weak incentive constraint for cost sharing rules. Despite the attention that the Aumann-Shapley rule received, it fails to meet this standard. To see this, consider the following:



Cost Sharing in Production Economies, Fig. 11 Paths for μ^{MS} , μ^{AS} , μ^{FM} induced by $q = (10; 10)$



Cost Sharing in Production Economies, Fig. 12 Paths for μ^{MS} , μ^{AS} , μ^{FM} induced by $q = (5; 10)$

$$c(y) = \frac{y_1 y_2}{y_1 + y_2}, \quad \text{and} \quad \mu_1^{\text{AS}}(x, c) = \frac{x_1 x_2^2}{(x_1 + x_2)^2}.$$

Then the latter expression is not monotonic in x_1 . One may show that the combinations of properties in Theorem 7.14 are incompatible with demand monotonicity. Now what kind of rules is demand monotonic? The classification of all such rules is too complex. We will restrict our attention to the additive rules with the *dummy* property, which comprises the idea that a player pays nothing if her good is for free:

Dummy If $\partial_i c(y) = 0$ for all y , then $\mu_i(x, c) = 0$ for all cost sharing problems $(x, c) \in P^n$.

Theorem 7.17 *Friedman (2004)*

- A cost sharing rule μ satisfies *dummy*, *additivity*, and *demand monotonicity* if and only if it is an (infinite) convex combination of rules generated by fixed paths which do not depend on the cost structure.
- A cost sharing rule μ satisfies *dummy*, *additivity*, and *scale invariance* if and only if it is an (infinite) convex combination of rules generated by scale invariant fixed paths which do not depend on the cost structure.

This theorem has some important implications. The Friedman-Moulin rule is the unique serial extension with the properties additivity, dummy, and demand monotonicity. As we mentioned before, μ^{FM} is not scale invariant. The only cost sharing rules satisfying all four of the above properties are *random order values*, i.e., a convex combination of marginal vectors of the stand-alone cost game (Weber 1988). μ^{SS} is the special element in this class of rules, giving equal weight to each marginal vector. Consider the following weak fairness property:

Equal Treatment Consider $(x, c) \in P^n$. If $c_x(S \cup \{i\}) = c_x(S \cup \{j\})$ for all i, j and $S \subset N \setminus \{i, j\}$, then $\mu_i(x, c) = \mu_j(x, c)$.

Within the class of random order values, μ^{SS} is the unique cost sharing rule satisfying equal treatment.

Strategic Properties of Fixed-Path Rules

Friedman (2002) shows that the fixed-path cost sharing rules essentially have the same strategic properties as the serial rule. The crucial result in this respect is the following.

Theorem 7.18 *Friedman (2002)*

Consider the demand game $G(\mu, c)$ where μ is a fixed-path cost sharing rule and $c \in C^n$ has strictly increasing partial derivatives. Then the corresponding set O^∞ of action profiles surviving the successive elimination of overwhelmed actions consists of a unique element.

As a corollary one may prove that the action profile in O^∞ is actually the unique Nash equilibrium of the game $G(\mu, c)$ and that it is strong as well. Moreover, Friedman (2002) shows that this Nash equilibrium can be reached through some learning dynamics. Then this means that the demand games induced by μ^{FM} and μ^{MS} have strong strategic properties. Notice that the above theorem is only one-way. There are other cost sharing rules, like the *axial serial rule*, having the same strategic characteristics.

Stochastic Cost Sharing Models

Throughout the previous sections, it was assumed that costs for service are known and in particular *deterministic*. This assumption seems way too strong for real-life situations, where the agents are participants in uncertain projects which may or may not be successful upon realization. Uncertainty may arise as (a) final service level is uncertain, or if (b) the cost involving service is not clear. As an example of the first source of uncertainty, consider large ICT projects. Complexity frequently causes uncertainty about the promised service levels for agents as applications may be dysfunctional or not as efficient as foreseen. The second source of uncertainty may also arise in

large projects, where time alone and changing prices may lead to uncertain outcomes and costs. Below two models will be discussed, where both sources of uncertainty are treated differently.

In practice agents commit themselves to a cost sharing rule *ex ante*, before the uncertainty is revealed. Then *ex post* there may be asymmetries in service levels for agents and/or differences in liabilities toward total costs. The two models show that *fairness* allows for *ex ante* and *ex post* interpretations. In the former sections, it was clear that pinning down *fairness* of a solution can already be an ambiguous task if only agents share heterogeneous characteristics which do not allow for a straightforward comparison in terms of a complete or partial ordering. This also holds for the combination of heterogeneous agents in those models with uncertainty.

Sharing Cost of Success and Failure of Projects

The first type of uncertainty in cost sharing problems is addressed in (Hougaard and Moulin 2017), which draws on ideas for the deterministic case in (Hougaard and Moulin 2014). Each agent has a binary inelastic demand for service of certain projects. *Ex ante* it is not clear which of the projects will actually work successfully and can be used to satisfy the agents' demand. The agents face the problem of sharing the cost of these unreliable (discrete) projects. *Ex post* it is clear what projects or parts of projects are functioning and which agents get full, partial, or no service. Two cost sharing rules are presented which take different positions regarding *ex ante* and *ex post* properties. This will be illustrated using the simple but illustrative example as in of the authors:

Example 8.1 Consider the following cost sharing problem among agents *Ann* and *Bob*: *Ann* is satisfied with service of only one of the projects *a* or *b*, denoted $D^A = \{a \vee b\}$, whereas *Bob* needs both projects to succeed, i.e., $D^B = \{a \wedge b\}$. Suppose that each project is of cost 1 and that probability that the probabilities that the projects are successful are independent and equal to p . In particular, the probability that *Bob* gets the required service equals p^2 . A pure *ex ante* idea of sharing costs would be to share the cost proportional to the probabilities of

getting service, that is, proportional to $p(2 - p)$ for *Ann* and p^2 for *Bob*. Then according to this cost sharing rule, cost shares are $2 - p$ for *Ann* and p for *Bob*. In particular, this means that *Ann* and *Bob* are considered equally liable toward both projects. A more subtle rule is the following *ex post* rule, which considers cost shares in relation to what is actually realized successfully. The *ex post* cost sharing rule proposes to share the costs equally if both *Ann* and *Bob* are served or none is served, and *Ann* pays if she is served but not *Bob*. This rule takes the expectation over the deterministic cost sharing problems after realization turns out to be success or failure (and this is the property of a rule that will be referred to as *independence of timing*). Then the corresponding cost shares are:

$$\begin{aligned} (\mu_A^{XP}, \mu_B^{XP}) &= (p^2 + (1 - p^2))\left(\frac{1}{2}, \frac{1}{2}\right) + 2p(1 - p)(1, 0) \\ &= \left(\frac{1}{2} + p(1 - p), \frac{1}{2} - p(1 - p)\right) \end{aligned}$$

The *needs priority* cost sharing rule determines for each realized combination of projects per project which of the agents' demands are fulfilled iff the project is successful and charges each of these agents equally for its cost. For instance, in the situation at hand, after successful realization of only project *a* only for *Ann*, the outcome is critical (*Bob* is not served anyways, since *b* has failed), and she is charged with the full cost 1. In the same fashion, *Bob* is charged for the full cost when both projects *a* and *b* are up and running, since each is critical to *Bob*, but not to *Ann*. If no project is a success, then costs are split equally. Then this leads to the following expected cost shares:

$$\begin{aligned} (\mu_A^{NP}, \mu_B^{NP}) &= \mathbb{P}[\neg a \wedge \neg b](1, 1) \\ &\quad + (\mathbb{P}[\neg a \wedge b] + \mathbb{P}[a \wedge \neg b])(2, 0) \\ &\quad + \mathbb{P}[a \wedge b](0, 2) = (1 - p)^2(1, 1) \\ &\quad + 2p(1 - p)(1, 0) + p^2(0, 2) \\ &= (1 + 2p - 3p^2, 1 - 2p + 3p^2) \end{aligned}$$

Notice that *Ann* and *Bob* are charged the same costs in case $p = \frac{2}{3}$. $\triangleleft$

The model and the above two cost sharing rules will now be formally defined. Let A be a set of risky and costly projects and $c(a)$ is the cost of

carrying out project a . A subset $X \subset A$ of projects is realized successfully with probability $p(X)$. Now the set of agents N has to share the aggregate cost $c(A) := \sum_{a \in A} c(a)$, no matter which of the projects are successfully running.

A *cost allocation problem under risk* for agents set N is an ordered triple (Q, p, c) where $Q = (N, A, D)$ and where $D = \{D^i\}_{i \in N}$ is the profile of (heterogeneous) demands. Here $D^i \subseteq 2^A$ is the *service set* of i , meaning that i is served iff projects in D^i are successfully carried out. The ex post cost sharing rule μ^{XP} calculates equal cost shares for the agents served at X , i.e., $S(Q; X) = \{i \in N \mid X \in D^i\}$, and determines the expected cost shares with respect to p :

$$\mu^{\text{XP}}(Q, p, c) = \left(\sum_X p(X) \cdot e[S(Q, X)] \right) c(A).$$

Here $E[S]$ stands for the vector in the unit simplex $\Delta(N)$ over N such that $E(S)_i = 1/|S|$ if $i \in S$ and 0 otherwise; in addition we put $E[\emptyset]_i = E[N]_i = 1/|N|$ for all $i \in N$.

Using additional notation $e[S^1, S^2] := e[S^1]$ if $S^1 \neq \emptyset$ and $e[\emptyset; S^2] = e[S^2]$, the *needs priority* cost sharing rule μ^{NP} is defined as

$$\mu^{\text{NP}}(Q, p, c) = \sum_{a \in A} \left(\sum_{X \subseteq P(A)} p(X) \cdot e[S(Q; X) \setminus S(Q; X \setminus a); S(Q; X)] \right) c(a).$$

It should immediately be clear that a major difference between μ^{XP} and μ^{NP} is that the latter exploits the full structure of which projects are critical for serving an agent upon realization or not. The result of Hougaard and Moulin (2017) is that a couple of structural properties (like the counterparts of additivity and separability in the former section) combined with rather weak fairness criteria single out the two cost sharing rules μ^{XP} and μ^{NP} .

Theorem 8.2 (Hougaard and Moulin 2017) *Among the cost sharing rules satisfying Cost Additivity, Independence of Timing, and Separable Across Projects*

- (i) *The ex post service rule is characterized by No Charge for No Service, Liable for Flexibility, and Liable for Single-Minded Needs*

- (ii) *The needs priority rule is characterized by No Charge for No Service and Useless is Free.*

Here *Liable for Flexibility* states that when an agent i 's needs are easier to fulfill than j 's, and i is therefore more likely to be served, then i is considered (weakly) more liable than j . *Liable for Single-Minded Needs* demands that those agents having demand for a single project are assigned the highest liabilities to that project. In contrast, *Useless is Free* states that the liability of an agent i toward a project α is 0 whenever there is an agent j being *single-minded* about α .

Sharing a Random Cost Under Limited Liabilities

The setting of Koster and Boonen (2019) is that of a multi-divisional firm with a central service unit – to which all divisions have equal access. Running this shared facility is costly, and the divisions are charged for the full and uncertain cost. So the distinction with the previous model in terms of uncertainty is that demands are all the same, but the only stochastic variable is the cost itself. Again we seek to share the ex post cost by cost sharing rules to which the agent commit themselves, before realization of the project. New element here is that the benevolent board of the firm puts bounds on the maximal liability of a division, subject to a feasible allocation of costs. The maximal liabilities of the divisions may differ due to risk capital allocations within the firm that limit the capacity to bear risk for the divisions (see, e.g., Kamiya and Zanjani 2017; Myers and Read 2001; Zanjani 2002).

Example 8.3 For instance, consider a three-agent cost sharing problem concerning a project with a random $C \sim U_n(0, 10)$. Suppose that for the three agents the vector of liabilities is given, i.e., $L = (2, 3, 8)$. This means that agent 1, 2, and 3, respectively, cannot be charged more than 2, 3, and 8, respectively. So these liabilities are high enough such that even the possible realization of the maximum cost $C_{\max} = 10$ can be shared, but not in a strict egalitarian way for each realization. $\triangleleft$

Constrained Stochastic Cost Sharing Problem

The cost of the project, denoted by C , is a bounded, nonnegative random variable on a fixed probability space (in the set $\mathcal{L}^\infty$). A *constrained stochastic cost sharing problem* among agents N is now an ordered pair $(L, C) \in \mathcal{L}^\infty \times \mathbb{R}_{++}^N$ where $L = (L_i)_{i \in N}$ is the profile of maximal liabilities of the agents. Costs are bounded, i.e., $C \leq C_{\max}$ for some $C_{\max}$, and it is assumed that the collective of agents can actually pay $C_{\max}$, or $\sum_i L_i > C_{\max}$. A cost sharing rule μ describes a profile of cost shares $\mu(L, C) \in \mathbb{R}^N$ such that (1) $\sum_{i \in N} \mu_i(L, C) = C$ (realized costs are always allocated) and (2) $\mu(L, C) \leq L$ (the allocation is feasible for constrained liabilities). Note that now we do not require nonnegativity of the cost shares.

Utility The objective of each agent is to minimize $V(X) = \mathbb{E}_{F^X}[X]$, where $X \in \mathcal{L}^\infty$ is interpreted as a future cost and F^X is a probability measure that may depend on the ordering of X . V is a *coherent* risk measure as in Artzner et al. (1999). An example of such coherent risk measure is the expected value, but also all dual utility functions as in Yaari (1987). This allows for a much wider spectrum of utilities, especially those reflecting risk averse agents.

An ex ante solution ψ assigns to each tuple (L, C, V) a set of cost shares, which now may depend on V as well – as opposed to a cost sharing rule.

An allocation of costs reconciles the possible asymmetric characteristics of the divisions upon realization of the costs. Ideally, costs are shared equally by the divisions for any realization, if the liabilities allow so. But in case costs are high, it may be the case that firms with higher liabilities should contribute more than those with lower liabilities. So basically fairness is restricted by liabilities of the divisions. It is shown that there is a cost sharing rule that treats the divisions fairly, in the sense that if each is endowed with the same coherent risk measure V , the divisions are set out to the same level of risk, which is encapsulated by the following property:

Egalitarianism

Solution ψ is V -egalitarian if $V(\psi_i(L, C, V)) = V(\psi_j(L, C, V))$ for all $i, j \in N$.

Local Symmetry As long as agents are not tight with respect to their liabilities, marginal incremental costs are shared equally:

$$\begin{aligned} \psi_i(c, L, V) < L_i, \psi_j(c, L, V) < L_j &\Rightarrow \frac{\partial}{\partial c} \psi_i(c, L, V) \\ &= \frac{\partial}{\partial c} \psi_j(c, L, V). \end{aligned}$$

Basically LS is a strong property by which a two-part tariff is characterized.

Theorem 8.4 (Koster and Boonen 2019) *Let ψ be a solution with the local symmetry property. Then, for all (L, C, V) , it holds $\psi(L, C, V) = \psi(L, 0, V) + r^{\text{UG}}(L - \psi(L, 0, V), C)$.*

The structure of the rule ψ is simple, as it determines a vector of transfers allocated to the divisions in absence of costs, and the uniform “gains” rationing method r^{UG} is applied to the remaining variable component of the costs. Importantly there is a unique vector of so-called transfers $t \in \mathbb{R}^N$ such that

$$\psi(L, C, V) = t + r^{\text{UG}}(L - t, C)$$

The transfers can be interpreted as correction factors for differences in the maximal liability of paying a fair share of contingent high cost levels. The fact that any solution with the property LS decomposes in this fashion pins down the solution once transfers are selected. To select transfers, we next focus on ex ante allocation rules ψ that are V -egalitarian. Say that ψ satisfies the property LS if LS holds for any fixed V . Then there is a unique cost sharing rule with the property LS that is V -egalitarian.

Theorem 8.5 (Koster and Boonen (2019)) *Let V a comonotonic risk measure and stochastic cost sharing problem (L, C) a constrained stochastic cost sharing problem with $C_{\max} < L(N)$ and $L > \frac{1}{n}V(C)$. Then, there is a unique V -egalitarian rule ψ which satisfies LS, and that is $\psi = \psi^E$ which is defined by*

$$\begin{aligned}\psi(L, C, V) &= \psi^E(L, C, V) := t \\ &\quad + r^{\text{UG}}(L - t, C),\end{aligned}\quad (35)$$

where t is a unique vector of transfers such that ψ is V -egalitarian.

An axiomatic characterization is discussed in Koster and Boonen (2019) and it uses a parallel characterization of r^{UG} as in Yeh (2008).

Example 8.6 Consider again Example 8.3. Let $V(C) = \mathbb{E}[C]$, which means that $V(C) = 5$. It is easily verified that the problem (L, C, V) satisfies admissibility (or $L(N) > C_{\max} = 10$) and V -sufficiency, i.e., $L_i > \frac{1}{3}V(C)$. Then, we derive a unique vector of transfers $t \approx (0.51, -0.13, -0.38)$ such that solution ψ^E is V -egalitarian. This solution ψ^E is displayed in Fig. 13. The solution is linear in the sense that marginal contributions due to cost increase are equally shared among the agents whose liabilities are not fully attained (Fig. 13). $\triangleleft$

Remark It is shown in Koster and Boonen (2019) that the rule can easily be adapted in favor of a benevolent planner with a more skewed distribution of risk.


Future Directions

So far, a couple of standard stylized models have been launched providing a theoretical basis for defining and studying cost sharing principles at a basic level. The list of references below indicates that this field of research is full in swing, both in theoretical and applied directions. Although it is hard to make a guess where developments lead to, a couple of future directions will be highlighted.

Informational Issues

The cost sharing problems we face in practice are shaped by unsure events. Costs itself maybe unsure, agents may have fluctuating and (un)predictable characteristics, and even the relevant agent set may give rise to uncertainty. Especially, this creates informational problems in the context of say demand games, where the theoretical foundation requires the agents to have full knowledge of the situation at hand. There are not so many papers capturing a more general situation, and Kolpin and Wilbur (2005) that was discussed earlier serves as a rare exception.

The presented models assume the information of costs for every contingent demand profile. Certainly within the continuous framework, this seems too much to ask for. Retrieving the necessary information is hindered not only by technical constraints but leads to new costs as well.



Cost Sharing in Production Economies, Fig. 13 Graphical illustration of the V -egalitarian solution $\psi(L, C, V) = t + r^{\text{UG}}(L - t, C)$ corresponding to Example 8.6. The dotted line represents $\psi_1^E(L, \cdot, V)$, the dashed line

$\psi_2^E(L, \cdot, V)$, and the solid line $\psi_3^E(L, \cdot, V)$. Here, the cutoff points $\gamma \approx (4.47, 7.75, 10)$ are the points where one of the agents' liabilities is met

Hougaard and Tind (2007) discusses data envelopment in cost sharing problems. A stochastic framework will be useful to study such *estimated* cost sharing problems. Other work focusing on informational coherence in cost sharing problems is Sprumont (2000). Related work is Albizuri et al. (2003), discussing mixtures of discrete and continuous cost sharing problems.

Budget Balance

In this overview, the proposed mechanisms are based on cost *sharing* rules. Another stream in implementation theory – at the other extreme of the spectrum – deals with cost *allocation* rules with no restrictions on the budget. Moulin (2010) compares the size of budget deficits relative to the overall efficiency of a mechanism.

Performance

Recall the performance indices measuring the welfare impact of different cost sharing rules. Moulin (2008) focuses on the continuous homogeneous production situations, with cost functions of specific types. There is still a need for a more general theory. In particular this could prove to be indispensable for analyzing the quality of cost sharing rules in a broader setup, the heterogeneous and Bayesian cost sharing problems.

Nonlinear Cost Sharing Rules

Most of the axiomatic literature is devoted to the analysis of cost sharing rules as linear operators. The additivity property is usually motivated as an accounting convention, but it serves merely as a tool by which some mathematical representation theorems apply. Besides the practical motivation, it is void of any ethical content. As Moulin (2002) underlines, there are hardly results on non-additive cost sharing rules – one of the reasons is that the mathematical analysis becomes notoriously hard. But – as a growing number of authors acknowledge – the usefulness of these mathematical techniques alone cannot justify the contribution of the property.

Acknowledgments I thank Hervé Moulin for his contribution as a referee. His valuable comments sharpened the formulation of the main ideas in this exposition.

Bibliography

- Aadland D, Kolpin V (1998) Shared irrigation cost: an empirical and axiomatic analysis. *Math Soc Sci* 35:203–218
- Aadland D, Kolpin V (2004) Environmental determinants of cost sharing. *J Econ Behav Organ* 53:495–511
- Albizuri MJ, Zarzuelo JM (2007) The dual serial cost-sharing rule. *Math Soc Sci* 53:150–163
- Albizuri MJ, Santos JC, Zarzuelo JM (2003) On the serial cost sharing rule. *Int J Game Theory* 31:437–446
- An M (1998) Logconcavity versus logconvexity, a complete characterization. *J Econ Theory* 80:350–369
- Archer A, Feigenbaum J, Krishnamurthy A, Sami R, Shenker S (2004) Approximation and collusion in multicast costsharing. *Games Econ Behav* 47:36–71
- Arin J, Iñarra E (2001) Egalitarian solutions in the core. *Int J Game Theory* 30:187–193
- Artzner P, Delbaen F, Eber J-M, Heath D (1999) Coherent measures of risk. *Math Financ* 9:203–228
- Atkinson AB (1970) On the measurement of inequality. *J Econ Theory* 2:244–263
- Aumann RJ (1959) Acceptable points in general cooperative n-person games. In: *Contributions to the theory of games*, vol IV. Princeton University Press, Princeton
- Aumann RJ, Maschler M (1985) Game theoretic analysis of a bankruptcy problem from the Talmud. *J Econ Theory* 36:195–213
- Aumann RJ, Shapley LS (1974) *Values of non-atomic games*. Princeton University Press, Princeton
- Baumol W, Bradford D (1970) Optimal departure from marginal cost pricing. *Am Econ Rev* 60:265–283
- Baumol W, Panzar J, Willig R (1988) *Contestable markets and the theory of industry structure*. Harcourt College Pub, Revised Edition
- Bergantino A, Coppejans L (1997) A game theoretic approach to the allocation of joint costs in a maritime environment: a case study. *Occasional papers* 44. Department of Maritime Studies and International Transport, University of Wales, Cardiff
- Billera LJ, Heath DC (1982) Allocation of shared costs: a set of axioms yielding a unique procedure. *Math Oper Res* 7:32–39
- Billera LJ, Heath DC, Raanan J (1978) Internal telephone billing rates: a novel application of non-atomic game theory. *Oper Res* 26:956–965
- Binmore K (2007) *Playing for real: a text on game theory*. Oxford University Press, Oxford
- Bird CG (1976) On cost allocation for a spanning tree: a game theoretic approach. *Networks* 6:335–350
- Bjorndal E, Hamers H, Koster M (2004) Cost allocation in a bank ATM network. *Math Meth Oper Res* 59:405–418
- Bochet O, Klaus B (2007) A note on Dasgupta, Hammond, and Maskin's (1979) domain richness condition. Discussion paper RM/07/039, ME-TEOR, Maastricht
- Bondareva ON (1963) Some applications of linear programming to the theory of cooperative games. *Probl Kybern* 10:119–139. [in Russian]

- Brânzei R, Ferrari G, Fragnelli V, Tijs S (2002) Two approaches to the problem of sharing delay costs in joint projects. *Ann Oper Res* 109:359–374
- Chen Y (2003) An experimental study of serial and average cost pricing mechanisms. *J Public Econ* 87:2305–2335
- Clarke EH (1971) Multipart pricing of public goods. *Public Choice* 11:17–33
- Dasgupta P, Hammond P, Maskin E (1979) The implementation of social choice rules: some general results on incentive compatibility. *Rev Econ Stud* 46:185–216
- Davis M, Maschler M (1965) The kernel of a cooperative game. *Nav Res Logist Q* 12:223–259
- de Frutos MA (1998) Decreasing serial cost sharing under economies of scale. *J Econ Theory* 79:245–275
- Demers A, Keshav S, Shenker S (1990) Analysis and simulation of a fair queueing algorithm. *J Internetworking* 1:3–26
- Denault M (2001) Coherent allocation of risk capital. *J Risk* 4:1
- Dewan S, Mendelson H (1990) User delay costs and internal pricing for a service facility. *Manag Sci* 36:1502–1517
- Dutta B, Ray D (1989) A concept of egalitarianism under participation constraints. *Econometrica* 57:615–635
- Dutta B, Ray D (1991) Constrained egalitarian allocations. *Games Econ Behav* 3:403–422
- Flam SD, Jourani A (2003) Strategic behavior and partial cost sharing. *Games Econ Behav* 43:44–56
- Fleurbay M, Sprumont Y (2009) Sharing the cost of a public good without subsidies. *J Public Econ Theory* 11:1–9
- Friedman E (2002) Strategic properties of heterogeneous serial cost sharing. *Math Soc Sci* 44:145–154
- Friedman E (2004) Paths and consistency in additive cost sharing. *Int J Game Theory* 32:501–518
- Friedman E, Moulin H (1999) Three methods to share joint costs or surplus. *J Econ Theory* 87:275–312
- Friedman E, Shenker S (1998) Learning and implementation on the Internet. Working paper 1998–21. Department of Economics, Rutgers University
- González-Rodríguez P, Herrero C (2004) Optimal sharing of surgical costs in the presence of queues. *Math Methods Oper Res* 59:435–446
- Granot D, Huberman G (1984) On the core and nucleolus of minimum cost spanning tree games. *Math Program* 29(1984):323–347
- Green J, Laffont JJ (1977) Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica* 45:427–438
- Groves T (1973) Incentives in teams. *Econometrica* 41:617–663
- Haimanko O (2000) Partially symmetric values. *Math Oper Res* 25:573–590
- Harsanyi J (1967) Games with incomplete information played by Bayesian players. *Manag Sci* 14:159–182
- Hart S, Mas-Colell A (1989) Potential, value, and consistency. *Econometrica* 57:589–614
- Haviv M (2001) The Aumann-Shapley price mechanism for allocating congestion costs. *Oper Res Lett* 29:211–215
- Henriet D, Moulin H (1996) Traffic-based cost allocation in a network. *RAND J Econ* 27:332–345
- Hougaard JL, Moulin H (2014) Sharing the cost of redundant projects. *Games Econ Behav* 87:339–352
- Hougaard JL, Moulin H (2017) Sharing the cost of risky projects. *Economic Theory*. <https://doi.org/10.1007/s00199-017-1034-3>
- Hougaard JL, Thorlund-Petersen L (2000) The stand-alone test and decreasing serial cost sharing. *Economic Theory* 16:355–362
- Hougaard JL, Thorlund-Petersen L (2001) Mixed serial cost sharing. *Math Soc Sci* 41:51–68
- Hougaard JL, Tind J (2007) Cost allocation and convex data envelopment. *Eur J Oper Res* 194:939–947
- Iñarra E, Isategui JM (1993) The Shapley value and average convex games. *Int J Game Theory* 22:13–29
- Israelsen D (1980) Collectives, communes, and incentives. *J Comp Econ* 4:99–124
- Jackson MO (2001) A crash course in implementation theory. *Soc Choice Welf* 18:655–708
- Joskow PL (1976) Contributions of the theory of marginal cost pricing. *Bell J Econ* 7:197–206
- Kalai E (1977) Proportional solutions to bargaining situations: interpersonal utility comparisons. *Econometrica* 45:1623–1630
- Kaminski M (2000) ‘Hydrolic’ rationing. *Math Soc Sci* 40:131–155
- Kamiya S, Zanjani G (2017) Egalitarian equivalent capital allocation. *N Am Actuar J* 21:382–396
- Kolpin V, Wilbur D (2005) Bayesian serial cost sharing. *Math Soc Sci* 49:201–220
- Koster M (2002) Concave and convex serial cost sharing. In: Borm P, Peters H (eds) *Chapters in game theory*. Kluwer, Dordrecht
- Koster M (2005) Sharing variable returns of cooperation. CeNDEF working paper 05-06. University of Amsterdam, Amsterdam
- Koster M (2006) Heterogeneous cost sharing, the directional serial rule. *Math Methods Oper Res* 64:429–444
- Koster M (2007) The Moulin-Shenker rule. *Soc Choice Welf* 29:271–293
- Koster M (2012) Consistent cost sharing. *Math Meth Oper Res* 75:1–28. <https://doi.org/10.1007/s00186-011-0372-3>
- Koster M, Boonen T (2019) Constrained stochastic cost allocation. *Math Soc Sc* 101:20–30
- Koster M, Tijs S, Borm P (1998) Serial cost sharing methods for multicommodity situations. *Math Soc Sci* 36:229–242
- Koster M, Molina E, Sprumont Y, Tijs ST (2002) Sharing the cost of a network: core and core allocations. *Int J Game Theory* 30:567–599
- Koster M, Reijnierse H, Voorneveld M (2003) Voluntary contributions to multiple public projects. *J Public Econ Theory* 5:25–50
- Koutsoupias E, Papadimitriou C (1999) Worst-case equilibria. In: 16th annual symposium on theoretical aspects of computer science, Trier, pp 404–413

- Legros P (1986) Allocating joint costs by means of the nucleolus. *Int J Game Theory* 15:109–119
- Leroux J (2004) Strategy-proofness and efficiency are incompatible in production economies. *Econ Lett* 85:335–340
- Leroux J (2008) Profit sharing in unique Nash equilibrium: characterization in the two-agent case. *Games Econ Behav* 62(2):558–572
- Littlechild SC, Owen G (1973) A simple expression for the Shapley value in a simple case. *Manag Sci* 20:370–372
- Littlechild SC, Thompson GF (1977) Aircraft landing fees: a game theory approach. *Bell J Econ* 8:186–204
- Maniquet F, Sprumont Y (1999) Efficient strategy-proof allocation functions in linear production economies. *Economic Theory* 14:583–595
- Maniquet F, Sprumont Y (2004) Fair production and allocation of an excludable nonrival good. *Econometrica* 72:627–640
- Maschler M (1990) Consistency. In: Ichiiishi T, Neyman A, Tauman Y (eds) *Game theory and applications*. Academic, New York, pp 183–186
- Maschler M (1992) The bargaining set, kernel and nucleolus. In: Aumann RJ, Hart S (eds) *Handbook of game theory with economic applications*, vol I. North-Holland, Amsterdam
- Maschler M, Reijniere H, Potters J (1996) Monotonicity properties of the nucleolus of standard tree games. *Int J Game Theory* 39:89–104
- Maskin E, Sjöström T (2002) Implementation theory. In: Arrow KJ, Sen AK, Suzumura K (eds) *Handbook of social choice and welfare*, vol I. North-Holland, Amsterdam
- Matsubayashi N, Umezawa M, Masuda Y, Nishino H (2005) Cost allocation problem arising in hub-spoke network systems. *Eur J Oper Res* 160:821–838
- McLean RP, Pazgal A, Sharkey WW (2004) Potential, consistency, and cost allocation prices. *Math Oper Res* 29:602–623
- Mirman L, Tauman Y (1982) Demand compatible equitable cost sharing prices. *Math Oper Res* 7:40–56
- Monderer D, Shapley LS (1996) Potential games. *Games Econ Behav* 14:124–143
- Moulin H (1987) Equal or proportional division of a surplus, and other methods. *Int J Game Theory* 16:161–186
- Moulin H (1994) Serial cost-sharing of an excludable public good. *Rev Econ Stud* 61:305–325
- Moulin H (1995a) Cooperative microeconomics: a game-theoretic introduction. Prentice Hall, London
- Moulin H (1995b) On additive methods to share joint costs. *Jpn Econ Rev* 46:303–332
- Moulin H (1996) Cost sharing under increasing returns: a comparison of simple mechanisms. *Games Econ Behav* 13:225–251
- Moulin H (1999) Incremental cost sharing: characterization by coalition strategy-proofness. *Soc Choice Welf* 16:279–320
- Moulin H (2000) Priority rules and other asymmetric rationing methods. *Econometrica* 68:643
- Moulin H (2002) Axiomatic cost and surplus-sharing. In: Arrow KJ, Sen AK, Suzumura K (eds) *Handbook of social choice and welfare*. Handbooks in economics, vol 19. North-Holland Elsevier, Amsterdam, pp 289–357
- Moulin H (2008) The price of anarchy of serial, average and incremental cost sharing. *Economic Theory* 36:379–405
- Moulin H (2010) An efficient and almost budget balanced cost sharing method. *Games Econ Behav* 70:107–131
- Moulin H, Shenker S (1992) Serial cost sharing. *Econometrica* 60:1009–1037
- Moulin H, Shenker S (1994) Average cost pricing versus serial cost sharing; an axiomatic comparison. *J Econ Theory* 64:178–201
- Moulin H, Shenker S (2001) Strategy-proof sharing of submodular cost: budget balance versus efficiency. *Economic Theory* 18:511–533
- Moulin H, Sprumont Y (2005) On demand responsiveness in additive cost sharing. *J Econ Theory* 125:1–35
- Moulin H, Sprumont Y (2006) Responsibility and cross-subsidization in cost sharing. *Games Econ Behav* 55:152–188
- Moulin H, Vohra R (2003) Characterization of additive cost sharing methods. *Econ Lett* 80:399–407
- Moulin H, Watts A (1997) Two versions of the tragedy of the commons. *Econ Des* 2:399–421
- Mutuswami S (2004) Strategyproof cost sharing of a binary good and the egalitarian solution. *Math Soc Sci* 48:271–280
- Myers SC, Read JA (2001) Capital allocation for insurance companies. *J Risk Insur* 68:545–580
- Myerson RR (1980) Conference structures and fair allocation rules. *Int J Game Theory* 9:169–182
- Myerson RR (1991) *Game theory: analysis of conflict*. Harvard University Press, Cambridge, MA
- Nash JF (1950) Equilibrium points in n-person games. *Proc Natl Acad Sci* 36:48–49
- O'Neill B (1982) A problem of rights arbitration from the Talmud. *Math Soc Sci* 2:345–371
- Osborne MJ (2004) *An introduction to game theory*. Oxford University Press, New York
- Osborne MJ, Rubinstein A (1994) *A course in game theory*. MIT Press, Cambridge
- Peleg B, Sudhölter P (2004) *Introduction to the theory of cooperative games*. Series C: theory and decision library series. Springer-Verlag Berlin Heidelberg
- Pérez-Castrillo D, Wettstein D (2006) An ordinal Shapley value for economic environments. *J Econ Theory* 127:296–308
- Potters J, Sudhölter P (1999) Airport problems and consistent allocation rules. *Math Soc Sci* 38:83–102
- Razzolini L, Reksulak M, Dorsey R (2004) An experimental evaluation of the serial cost sharing rule. *Theor Decis* 63:283–314
- Ritzberger K (2002) *Foundations of non-cooperative game theory*. Oxford University Press, Oxford
- Rosenthal RW (1973) A class of games possessing pure-strategy Nash equilibria. *J Econ Theory* 2:65–67

- Roth AE (1988) The Shapley value, essays in honor of Lloyd S. Shapley. Cambridge University Press, Cambridge, pp 307–319
- Samet D, Tauman Y, Zang I (1984) An application of the Aumann-Shapley prices for cost allocation in transportation problems. *Math Oper Res* 9:25–42
- Sánchez SF (1997) Balanced contributions axiom in the solution of cooperative games. *Games Econ Behav* 20:161–168
- Sandmark M (1999) Production games under uncertainty. *Comput Econ* 14:237–253
- Schmeidler D (1969) The nucleolus of a characteristic function game. *SIAM J Appl Math* 17:1163–1170
- Shapley LS (1953) A value for n -person games. *Ann Math Study* 28:307–317. Princeton University Press, Princeton
- Shapley LS (1967) On balanced sets and cores. *Nav Res Logist Q* 14:453–460
- Shapley LS (1969) Utility comparison and the theory of games. In: La decision: Aggregation et dynamique des ordres de preference. Editions du Centre National de la Recherche Scientifique, Paris, pp 251–263. Also in Roth AE (ed) (1988) The Shapley value, essays in honor of Lloyd S. Shapley. Cambridge University Press, Cambridge, pp 307–319
- Shapley LS (1971) Cores of convex games. *Int J Game Theory* 1:1–26
- Sharkey W (1982) Suggestions for a game-theoretic approach to public utility pricing and cost allocation. *Bell J Econ* 13:57–68
- Sharkey W (1995) Network models in economics. In: Ball MO et al (eds) Network routing. Handbook in operations research and management science, vol 8. North-Holland, Amsterdam
- Shubik M (1962) Incentives, decentralized control, the assignment of joint cost, and internal pricing. *Manag Sci* 8:325–343
- Skorin-Kapov D (2001) On cost allocation in hub-like networks. *Ann Oper Res* 106:63–78
- Skorin-Kapov D, Skorin-Kapov J (2005) Threshold based discounting network: the cost allocation provided by the nucleolus. *Eur J Oper Res* 166:154–159
- Sprumont Y (1998) Ordinal cost sharing. *J Econ Theory* 81:126–162
- Sprumont Y (2000) Coherent cost sharing. *Games Econ Behav* 33:126–144
- Sprumont Y (2005) On the discrete version of the Aumann-Shapley cost-sharing method. *Econometrica* 73:1693–1712
- Sprumont Y, Ambec S (2002) Sharing a river. *J Econ Theory* 107:453–462
- Sprumont Y, Moulin H (2007) Fair allocation of production externalities: recent results, *Revue d'Économie Politique* 2007/1 (117)
- Sudhölter P (1998) Axiomatizations of game theoretical solutions for oneoutput cost sharing problems. *Games Econ Behav* 24:42–71
- Suijs J, Borm P, Hamers H, Koster M, Quant M (2005) Communication and cooperation in public network situations. *Ann Oper Res* 137:117–140
- Tauman Y (1988) The Aumann-Shapley prices: a survey. In: Roth A (ed) The Shapley value. Cambridge University Press, Cambridge, pp 279–304
- Thomas LC (1992) Dividing credit-card costs fairly. *IMA J Math Appl Bus Ind* 4:19–33
- Thomson W (1996) Consistent allocation rules. Mimeo, Economics Department, University of Rochester, Rochester
- Thomson W (2001) On the axiomatic method and its recent applications to game theory and resource allocation. *Soc Choice Welf* 18:327–386
- Tijs SH, Driessen TSH (1986) Game theory and cost allocation problems. *Manag Sci* 32:1015–1028
- Tijs SH, Koster M (1998) General aggregation of demand and cost sharing methods. *Ann Oper Res* 84:137–164
- Timmer J, Borm P, Tijs S (2003) On three Shapley-like solutions for cooperative games with random payoffs. *Int J Game Theory* 32:595–613
- van de Nouweland A, Tijs SH (1995) Cores and related solution concepts for multi-choice games. *Math Methods Oper Res* 41:289–311
- von Neumann J, Morgenstern O (1944) Theory of games and economic behavior. Princeton University Press, Princeton
- Watts A (2002) Uniqueness of equilibrium in cost sharing games. *J Math Econ* 37:47–70
- Weber RJ (1988) Probabilistic values for games. In: Roth AE (ed) The Shapley value. Cambridge University Press, Cambridge
- Yaari ME (1987) The dual theory of choice under risk. *Econometrica* 55:95–115
- Yeh CH (2008) Secured lower bound, composition up, and minimal rights first for bankruptcy problems. *J Math Econ* 44:925–932
- Young HP (1985a) Producer incentives in cost allocation. *Econometrica* 53:757–765
- Young HP (1985b) Monotonic solutions of cooperative games. *Int J Game Theory* 14:65–72
- Young HP (1985c) Cost allocation: methods, principles, applications. North-Holland, Amsterdam
- Young HP (1988) Distributive justice in taxation. *J Econ Theory* 44:321–335
- Young HP (1994) Cost allocation. In: Aumann RJ, Hart S (eds) Handbook of game theory, vol II. Elsevier, Amsterdam, pp 1193–1235
- Young HP (1998) Cost allocation, demand revelation, and core implementation. *Math Soc Sci* 36:213–229
- Zanjani G (2002) Pricing and capital allocation in catastrophe insurance. *J Financ Econ* 65:283–305

Market Games and Clubs

Myrna Wooders

Department of Economics, Vanderbilt University,
Nashville, TN, USA

Article Outline

Glossary

Definition of the Subject

Introduction

Transferable Utility Games; Some Standard
Definitions

A Market

Market-Game Equivalence

Equivalence of Markets and Games with Many
Players

Cores and Approximate Cores

Nonemptiness and Convergence of Approximate
Cores of Large Games

Shapley Values of Games with Many Players

Economies with Clubs

With a Continuum of Players

Other Related Concepts and Results

Some Remarks on Markets and More General
Classes of Economies

Conclusions and Future Directions

Bibliography

Glossary

An economy We use the term ‘economy’ to describe any economic setting, including economies with clubs, where the worth of club members may depend on the characteristics of members of the club, economies with pure public goods, local public goods (public goods subject to crowding and/or congestion), economies with production where what can be produced and the costs of production may depend on the characteristics of the individuals

involved in production, and so on. A *large economy* has many participants.

Asymptotic negligibility A pregame satisfies asymptotic negligibility if vanishingly small groups can have only negligible effects on per capita payoffs.

Club A club is a group of agents or players that forms for the purpose of carrying out some activity, such as providing a local public good.

Core The core of a game is the set (possibly empty) of feasible outcomes – divisions of the worths arising from coalition formation among the players of the game – that cannot be improved upon by any coalition of players.

Game A (cooperative) game (in characteristic form) is defined simply as a finite set of players and a function or correspondence ascribing a worth (a non-negative real number, interpreted as an idealized money) to each nonempty subset of players, called a group or coalition

Market games A market game is a game derived from a market. Given a market and a group of agents we can determine the total utility (measured in money) that the group can achieve using only the endowments belonging to the group members, thus determining a game.

Market A market is defined as a private goods economy in which all participants have utility functions that are linear in (at least) one commodity (money).

Payoff vector A payoff vector is a vector listing a payoff (an amount of utility or money) for each player in the game.

Per capita boundedness A pregame satisfies per capita boundedness if the supremum of the average worth of any possible group of players (the per capita payoff) is finite.

Pregame A pair, consisting of a set of player types (attributes or characteristics) and a function mapping finite lists of characteristics (repetitions allowed) into the real numbers. In interpretation, the pregame function ascribes a

worth to every possible finite group of players, where the worth of a group depends on the numbers of players with each characteristic in the group. A pregame is used to generate games with arbitrary numbers of players.

Price taking equilibrium A price taking equilibrium for a market is a set of prices, one for each commodity, and an allocation of commodities to agents so that each agent can afford his part of the allocation, given the value of his endowment.

Shapley value The Shapley value of a game is a feasible outcome of a game in which all players are assigned their expected marginal contribution to a coalition when all orders of coalition formation are equally likely.

Small group effectiveness A pregame satisfies small group effectiveness if almost all gains to collective activities can be realized by cooperation only within arbitrarily small groups (coalitions) of players.

Totally balanced game A game is totally balanced if the game and every subgame of the game (a game with player set taken as some subset of players of the initially given game) has a nonempty core.

Definition of the Subject

The equivalence of markets and games concerns the relationship between two sorts of structures that appear fundamentally different – markets and games. Shapley and Shubik (1969) demonstrates that: (1) games derived from markets with concave utility functions generate totally balanced games where the players in the game are the participants in the economy and (2) every totally balanced game generates a market with concave utility functions. A particular form of such a market is one where the commodities are the participants themselves, a labor market for example.

But markets are very special structures, more so when it is required that utility functions be concave. Participants may also get utility from belonging to groups, such as marriages, or clubs, or productive coalitions. It may be that participants in an economy even derive utility

(or disutility) from engaging in processes that lead to the eventual exchange of commodities. The question is when are such economic structures equivalent to markets with concave utility functions.

This paper summarizes research showing that a broad class of large economies generate balanced market games. The economies include, for example, economies with clubs where individuals may have memberships in multiple clubs, with indivisible commodities, with nonconvexities and with non-monotonicities. The main assumption are: (1) that an option open to any group of players is to break into smaller groups and realize the sum of the worths of these groups, that is, essential super-additivity is satisfied and: (2) relatively small groups of participants can realize almost all gains to coalition formation.

The equivalence of games with many players and markets with many participants indicates that relationships obtained for markets with concave utility functions and many participants will also hold for diverse social and economic situations with many players. These relationships include: (a) equivalence of the core and the set of competitive outcomes; (b) the Shapley value is contained in the core or approximate cores; (c) the equal treatment property holds – that is, both market equilibrium and the core treat similar players similarly. These results can be applied to diverse economic models to obtain the equivalence of cooperative outcomes and competitive, price taking outcomes in economies with many participants and indicate that such results hold in yet more generality.

Introduction

One of the subjects that has long intrigued economists and game theorists is the relationship between games, both cooperative and noncooperative, and economies. Seminal works making such relationships include Shubik (1959b), Debreu and Scarf (1963), Aumann (1964), Shapley and Shubik (1969, 1975) and Aumann and Shapley (1974), all connecting outcomes of price-taking behavior in large economies with

cores of games. See also Shapley and Shubik (1977) and an ongoing stream of papers connecting strategic behavior to market behavior. Our primary concern here, however, is not with the equivalence of outcomes of solution concepts for economies, as is Debreu and Scarf (1963) or Aumann and Dreze (1974) for example, but rather with equivalences of the *structures* of markets and games. Solution concepts play some role, however, in establishing these equivalences and in understanding the meaning of the equivalence of markets and games.

In this entry, following Shapley and Shubik (1969), we focus on markets in which utility functions of participants are quasi-linear, that is, the utility function u of a participant can be written as $u(x, \xi) = \hat{u}(x) + \xi$ where $x \in \mathbb{R}_+^L$ is a commodity bundle, $\xi \in \mathbb{R}$ is interpreted as money and $\hat{u}$ is a continuous function. Each participant in an economy has an endowment of commodities and, without any substantive loss of generality, it is assumed that no money is initially endowed. The price of money is assumed equal to one. A price taking equilibrium for a market then consists of a price vector $p \in \mathbb{R}^L$ for the commodities. And an assignment of commodities to participants such that: the total amounts of commodities assigned to participants equals the total amount of commodities with which participants are endowed and; given prices, each participant can afford his assignment of commodities and no participant, subject to his budget constraint, can afford a preferred commodity bundle.

We also treat games with side payments, alternatively called games with transferable utility or, in brief, TU games. Such a game consists of a finite set N of players and a worth function that assigns to each group of players $S \subset N$ a real number $v(S) \in \mathbb{R}_+$, called the worth of the group. In interpretation, $v(S)$ is the total payoff that a group of players can realize by cooperation. A central game-theoretic concept for the study of games is the core. The core consists of those divisions of the maximal total worth achievable by cooperation among the players in N so that each group of players is assigned at least its worth. A game is balanced if it has a nonempty core and totally balanced if all subgames of the

game have nonempty cores. A subgame of a game is simply a group of players $S \subset N$ and the worth function restricted to that group and the smaller groups that it contains.

Given a market any feasible assignment of commodities to the economic participants generates a total worth of each group of participants. The worth of a group of participants (viewed as players of a game) is the maximal total utility achievable by the members of the group by allocating the commodities they own among themselves. In this way a market generates a game – a set of players (the participants in the economy) and a worth for each group of players.

Shapley and Shubik (1969) demonstrate that any market where all participants have concave, monotonic increasing utility functions generates a totally balanced game and that any totally balanced game generates a market, thus establishing an equivalence between a class of markets and totally balanced cooperative games. A particular sort of market is canonical; one where each participant in the market is endowed with one unit of a commodity, his “type”. Intuitively, one might think of the market as one where each participant owns one unit of himself or of his labor.

In the last 20 years or so there has been substantial interest in broader classes of economies, including those with indivisibilities, non-monotonicities, local public goods or clubs, where the worth of a group depends not only on the private goods endowed to members of the group but also on the characteristics of the group members. For example, the success of the marriage of a man and a woman depends on their characteristics and on whether their characteristics are complementary. Similarly, the output of a machine and a worker using the machine depends on the quality and capabilities of the machine and how well the abilities of the worker fit with the characteristics of the machine – a concert pianist fits well with an high quality piano but perhaps not so well with a sewing machine. Or how well a research team functions depends not only on the members of the team but also on how well they interact. For simplicity, we shall refer to these economies as club economies. Such economies can be modeled as cooperative games.

In this entry we discuss and summarize literature showing that economies with many participants are approximated by markets where all participants have the same concave utility function and for which the core of the game is equivalent to the set of price-taking economic equilibrium payoffs. The research presented is primarily from Shubik and Wooders (1982a), Wooders (1997) and earlier papers due to this author. For the most recent results in this line of research we refer the reader to Wooders (2007, 2008a, b). We also discuss other related works throughout the course of the entry. The models and results are set in a broader context in the conclusions.

The importance of the equivalence of markets and games with many players relates to the hypothesis of perfect competition, that large numbers of participants leads to price-taking behavior, or behavior “as if” participants took prices as given. Von Neumann and Morgenstern perceived that even though individuals are unable to influence market prices and cannot benefit from strategic behavior in large markets, large “coalitions” might form. Von Neumann and Morgenstern write:

It is neither certain nor probable that a mere increase in the number of participants might lead *in fine* to the conditions of free competition. The classical definitions of free competition all involve further postulates besides this number. E.g., it is clear that if certain great groups of individuals will -for any reason whatsoever-act together, then the great number of participants may not become effective; the decisive exchanges may take place directly between large “coalitions”, few in number and not between individuals, many in number acting independently. . . . Any satisfactory theory . . . will have to explain when such big coalitions will or will not be formed – i.e., when the large numbers of participants will become effective and lead to more or less free competition.

The assumption that small groups of individuals cannot affect market aggregates, virtually taken for granted by von Neumann and Morgenstern, lies behind the answer to the question they pose. The results presented in this entry suggest that the great number of participants will become effective and lead to more or less free competition when *small* groups of participants cannot significantly affect market outcomes. Since all or almost all gains to collective activities can be captured by relatively small groups, large

groups gain no market power from size; in other words, large groups are inessential. That large groups are inessential is equivalent to small group effectiveness (Wooders 1992b). A remarkable feature of the results discussed in this essay is they are independent of any particular economic structure.

Transferable Utility Games; Some Standard Definitions

Let (N, v) be a pair consisting of a finite set N , called a *player set*, and a function v , called a *worth function*, from subsets of N to the real numbers $\mathbb{R}$ with $v(\emptyset) = 0$. The pair (N, v) is a *TU game* (also called a game with side payments). Nonempty subsets S of N are called *groups* (of players) and the number of members of the group S is given by $|S|$. Following is a simple example.

Example 1 A glove game: Suppose that we can partition a player set N into two groups, say N_1 and N_2 . In interpretation, a member of N_1 is endowed with a right-hand (RH) glove and a member of N_2 is endowed with a left-hand (LH) glove. The worth of a pair of gloves is \$1, and thus the worth of a group of players consisting of player $i \in N_1$ and player $j \in N_2$ is \$1. The worth of a single glove and hence of a one-player group is \$0. The worth of a group $S \subset N$ is given by $v(S) = \min \{|S \cap N_1|, |S \cap N_2|\}$. The pair (N, v) is a game.

A *payoff vector* for a game (N, v) is a vector $\bar{u} \in \mathbb{R}^N$. We regard vectors in finite dimensional Euclidean space $\mathbb{R}^T$ as functions from T to $\mathbb{R}$, and write $\bar{u}_i$ for the i th component of $\bar{u}$, etc. If $S \subset T$ and $\bar{u} \in \mathbb{R}^T$, we shall write $\bar{u}_S := (\bar{u}_i \mid i \in S)$ for the restriction of $\bar{u}$ to S . We write 1_S for the element of $\mathbb{R}^S$ all of whose coordinates are 1 (or simply 1 if no confusion can arise.) A payoff vector u is *feasible* for a group $S \subset N$ if

$$\bar{u}(S) \stackrel{\text{def}}{=} \sum_{i \in S} \bar{u}^i \leq \sum_{k=1}^K v(S^k) \quad (1)$$

for some partition $\{S^1, \dots, S^K\}$ of S .

Given $\varepsilon \geq 0$, a payoff vector $\bar{u} \in \mathbb{R}^N$ is in the *weak ε -core* of the game (N, v) if it is feasible and if there is a group of players $N^0 \subset N$ such that

$$\frac{N \setminus N^0}{|N|} \leq \varepsilon \quad (2)$$

and, for all groups $S \subset N^0$,

$$\bar{u}(S) \geq v(S) - \varepsilon|S| \quad (3)$$

where $|S|$ is the cardinality of the set S . (It would be possible to use two different values for epsilon in expressions (2) and (3). For simplicity, we have chosen to take the same value for epsilon in both expressions.)

A payoff vector $\bar{u}$ is in the *uniform ε -core* (or simply in the ε -core) if it is feasible and if (3) holds for *all* groups $S \subset N$. When $\varepsilon = 0$, then both notions of ε -cores will be called simply the *core*.

The glove game (N, v) described in Example 1 has the happy feature that the core is always nonempty. For the game to be of interest, we will suppose that there is least one player of each type (that is, there is at least one player with a RH glove and one player with a LH glove). If $|N_1| = |N_2|$ any payoff vector assigning the same share of a dollar to each player with a LH glove and the remaining share of a dollar to each player with a RH glove is in the core. If there are more players of one type, say $|N_1| > |N_2|$ for specificity, then any payoff vector in the core assigns \$1 to each player of the scarce type; that is, players with a RH glove each receive 0 while players with a LH glove each receive \$1.

Not all games have nonempty cores, as the following example illustrates.

Example 2 (A simple majority game with an empty core) Let $N = \{1, 2, 3\}$ and define the function v as follows:

$$v(S) = \begin{cases} 0 & \text{if } |S| = 1, \\ 1 & \text{otherwise.} \end{cases}$$

It is easy to see that the core of the game is empty. For if a payoff vector $\bar{u}$ were in the core,

then it must hold that for any $i \in N$, $\bar{u}_i \geq 0$ and for any $i, j \in N$, $\bar{u}_i + \bar{u}_j \geq 1$. Moreover, feasibility dictates that $\bar{u}_1 + \bar{u}_2 + \bar{u}_3 \leq 1$. This is impossible; thus, the core is empty.

Before leaving this example, let us ask whether it would be possible to subsidize the players by increasing the payoff to the total player set N and, by doing so, ensure that the core of the game with a subsidy is nonempty. We leave it to the reader to verify that if $v(N)$ were increased to $\$3/2$ (or more), the new game would have a nonempty core.

Let (N, v) be a game and let $i, j \in N$. Then players i and j are *substitutes* if, for all groups $S \subset N$ with $i, j \notin S$ it holds that

$$v(S \cup \{i\}) = v(S \cup \{j\}).$$

Let (N, v) be a game and let $\bar{u} \in \mathbb{R}^N$ be a payoff vector for the game. If for all players i and j who are substitutes it holds that $\bar{u}_i = \bar{u}_j$ then u has the *equal treatment property*. Note that if there is a partition of N into T subsets, say $N_1, \dots, N_T$ where all players in each subset N_t are substitutes for each other, then we can $\bar{u}$ by a vector $\bar{\bar{u}} \in \mathbb{R}^T$ where, for each t , it holds that $\bar{\bar{u}}_t = \bar{u}_i$ for all $i \in N_t$.

Essential Superadditivity

We wish to treat games where the worth of a group of players is independent of the total player set in which it is embedded and an option open to the members of a group is to partition themselves into smaller groups; that is, we treat games that are *essentially superadditive*. This is built into our the definition of feasibility above, (1). An alternative approach, which would still allow us to treat situations where it is optimal for players to form groups smaller than the total player set, would be to assume that v is the “superadditive cover” of some other worth function v' . Given a not-necessarily-superadditive function v' , for each group S define $v(S)$ by:

$$v(S) = \max \sum v'(S^k) \quad (4)$$

where the maximum is taken over all partitions $\{S^k\}$ of S ; the function v is the *superadditive cover*

of v' . Then the notion of feasibility requiring that a payoff vector $\bar{u}$ is feasible only if

$$\bar{u}(N) \leq v(N), \quad (5)$$

gives an equivalent set of feasible payoff vectors to those of the game (N, v') with the definition of feasibility given by (1).

The following Proposition may be well known and is easily proven. This result was already well understood in Gillies (1953) and applications have appeared in a number of papers in the theoretical literature of game theory; see, for example (for $\varepsilon = 0$) Aumann and Dreze (1974) and Kaneko and Wooders (1982). It is also well known in club theory and the theory of economies with many players and local public goods.

Proposition 1 Given $\varepsilon \geq 0$, let (N, v') be a game. A payoff vector $u \in R^N$ is in the weak, respectively uniform, ε -core of (N, v') if and only if it is in the weak, respectively uniform, ε -core of the superadditive cover game, say (N, v) , where v is defined by (4).

A Market

In this section we introduce the definition, from Shapley and Shubik (1969), of a market. Unlike Shapley and Shubik, however, we do not assume concavity of utility functions. A *market* is taken to be an economy where all participants have continuous utility functions over a finite set of commodities that are all linear in one commodity, thought of as an “idealized” money. Money can be consumed in any amount, possibly negative. For later convenience we will consider an economy where there is a finite set of types of participants in the economy and all participants of the same type have the same endowments and preferences.

Consider an economy with $T + 1$ types of commodities. Denote the set of participants by

$$N = \{(t, q) : t = 1, \dots, T, \text{ and } q = 1, \dots, n_t\}.$$

Assume that all participants of the same type, (t, q) , $q = 1, \dots, n_t$ have the same utility functions given by

$$\hat{u}_t(y, \xi) = u_t(y) + \xi$$

where $y \in \mathbb{R}_+^T$ and $\xi \in \mathbb{R}$. Let $a^{tq} \in \mathbb{R}_+^T$ be the *endowment* of the (t, q) th player of the first T commodities. The *total endowment* is given by $\sum_{(t,q)} a^{tq} \in \mathbb{R}_+^T$. For simplicity and without loss of generality, we can assume that no participant is endowed with any nonzero amount of the $(T + 1)^{\text{th}}$ good, the “money” or medium of exchange. One might think of utilities as being measured in money. It is because of the transferability of money that utilities are called “transferable”.

Remark 1 Instead of assuming that money can be consumed in negative amounts one might assume that endowments of money are sufficiently large so that no equilibrium allocates any participant a negative amount of money. For further discussion of transferable utility see, for example, Bergstrom and Varian (1985) or Kaneko and Wooders (2004).

Given a group $S \subset N$, a *S-allocation of commodities* is a set

$$\left\{ \{y^{tq}, \xi^{tq}\} \in \mathbb{R}_+^T \times \mathbb{R} : \sum_{(t,q) \in S} y^{tq} \leq \sum_{(t,q) \in S} a^{tq} \text{ and } \sum_{(t,q) \in S} \xi^{tq} \leq 0 \right\}$$

that is, a *S-allocation* is a redistribution of the commodities owned by the members of S among themselves and monetary transfers adding up to no more than zero. When $S = N$, a *S-allocation* is called simply an *allocation*. With the price of the $(T + 1)^{\text{th}}$ commodity ξ set equal to 1, a *competitive outcome* is a price vector p in R^T , listing prices for the first T commodities, and an allocation $\{(y^{tq}, \xi^{tq}) \in R^T \times R : (t, q) \in N\}$ for which

$$\begin{aligned} & (a) \quad u_t(y^{tq}) - p \cdot (y^{tq} - a^{tq}) \geq u_t(\hat{y}) \\ & \quad - p \cdot (\hat{y} - a^{tq}) \text{ for all } \hat{y} \in R_+^T, (t, q) \in N. \\ & (b) \quad \sum_{(t,q) \in N} y^{tq} = \sum_{(t,q)} a^{tq} = \bar{y} \\ & (c) \quad \xi^{tq} = p \cdot (y^{tq} - a^{tq}) \text{ for all } (t, q) \in N \text{ and} \\ & (d) \quad \sum_{(t,q) \in N} \xi^{tq} = 0. \end{aligned} \quad (6)$$

Given a competitive outcome with allocation $\{(y^{tq}, \xi^{tq}) \in R_+^T \times R : (t, q) \in N\}$ and price vector p ,

the competitive payoff to the $(t, q)^{th}$ participant is $u(y^{tq}) - p \cdot (y^{tq} - a^{tq})$. A competitive payoff vector is given by

$$(u(y^{tq}) - p \cdot (y^{tq} - a^{tq}) : (t, q) \in N).$$

In the following we will assume that for each t , all participants of type t have the same endowment; that is, for each t , it holds that $a^{tq} = a^{tq'}$ for all $q, q' = 1, \dots, n_t$. In this case, every competitive payoff has the equal treatment property;

$$u_t(y^{tq}) - p \cdot (y^{tq} - a^{tq}) = u_t(y^{tq'}) - p \cdot (y^{tq'} - a^{tq'})$$

for all q, q' and for each t . It follows that a competitive payoff vector can be represented by a vector in $\mathbb{R}^T$ with one component for each player type.

It is easy to generate a game from the data of an economy. For each group of participants $S \subset N$, define

$$v(S) = \max_{tq \in S} \sum_{tq \in S} u_t(y^{tq}, \xi^{tq})$$

where the maximum is taken over the set of S -allocations. Let (N, v) denote a game derived from a market.

Under the assumption of concavity of the utility functions of the participants in an economy, Shapley and Shubik (1969) show that a competitive outcome for the market exists and that the competitive payoff vectors are in the core of the game. (Since (Debreu and Scarf 1963), such results have been obtained in substantially more general models of economies.)

Market-Game Equivalence

To facilitate exposition of the theory of games with many players and the equivalence of markets and games, we consider games derived from a common underlying structure and with a fixed number of types of players, where all players of the same type are substitutes for each other.

Pregames

Let T be a positive integer, to be interpreted as a number of player types. A profile $s = (s_1, \dots, s_T) \in \mathbf{Z}_+^T$, where $\mathbf{Z}_+^T$ is the T -fold Cartesian product of the non-negative integers $\mathbf{Z}_+$, describes a group of players by the numbers of players of each type in the group. Given profile s , define the *norm* or *size* of s by

$$\|s\| \stackrel{\text{def}}{=} \sum_t s_t,$$

simply the total number of players in a group of players described by s . A *subprofile* of a profile $n \in \mathbf{Z}_+^T$ is a profile s satisfying $s \leq n$. A *partition* of a profile s is a collection of subprofiles $\{s^k\}$ of n , not all necessarily distinct, satisfying

$$\sum_k s^k = s.$$

A partition of a profile is analogous to a partition of a set except that all members of a partition of a set are distinct.

Let Ψ be a function from the set of profiles $\mathbf{Z}_+^T$ to $\mathbb{R}_+$ with $\Psi(0) = 0$. The value $\Psi(s)$ is interpreted as the total payoff a group of players with profile s can achieve from collective activities of the group membership and is called the *worth of the profile* s .

Given Ψ , define a worth function Ψ^* , called the *superadditive cover* of Ψ , by

$$\Psi^*(s) \stackrel{\text{def}}{=} \max_k \sum_k \Psi(s^k),$$

where the maximum is taken over the set of all partitions $\{s^k\}$ of s . The function Ψ is said to be *superadditive* if the worth functions Ψ and Ψ^* are equal.

We define a *pregame* as a pair (T, Ψ) where $\Psi: \mathbf{Z}_+^T \rightarrow \mathbb{R}_+$. As we will now discuss, a pregame can be used to generate multiple games. To generate a game from a pregame, it is only required to specify a total player set N and the numbers of players of each of T types in the set. Then the pregame can be used to assign a worth to every group of players contained in the total player set, thus creating a game.

A game determined by the pregame (T, Ψ) , which we will typically call a *game* or a *game with side payments*, is a pair $[n; (T, \Psi)]$ where n is a profile. A *subgame* of a game $[n; (T, \Psi)]$ is a pair $[s; (T, \Psi)]$ where s is a subprofile of n .

With any game $[n; (T, \Psi)]$ we can associate a game (N, v) in the form introduced earlier as follows: Let

$$N = \{(t, q) : t = 1, \dots, T \text{ and } q = 1, \dots, n_t\}$$

be a *player set* for the game. For each subset $S \subset N$ define the *profile* of S , denoted by $\text{prof}(S) \in \mathbf{Z}_+^T$, by its components

$$\text{prof}(s)_T \stackrel{\text{def}}{=} |\{S \cap \{t', q\} : t' = t \text{ and } q = 1, \dots, n_t\}|$$

and define

$$v(s) \stackrel{\text{def}}{=} \Psi(\text{prof}(S)).$$

Then the pair (N, v) satisfies the usual definition of a game with side payments. For any $S \subset N$, define

$$v^*(s) \stackrel{\text{def}}{=} \Psi^*(\text{prof}(S)).$$

The game (N, v^*) is the *superadditive cover* of (N, v) .

A *payoff vector* for a game (N, v) is a vector $\bar{u} \in \mathbf{R}^N$. For each nonempty subset S of N define

$$\bar{u}(S) \stackrel{\text{def}}{=} \sum_{(t, q) \in S} \bar{u}^{tq}.$$

A payoff vector $\bar{u}$ is *feasible* for S if

$$\bar{u}(S) \leq v^*(S) = \Psi^*(\text{prof}(S)).$$

If $S = N$ we simply say that the payoff vector $\bar{u}$ is *feasible* if

$$\bar{u}(N) \leq v^*(N) = \Psi^*(\text{prof}(N)).$$

Note that our definition of feasibility is consistent with essential superadditivity; a group can realize at least as large a total payoff as it can

achieve in any partition of the group and one way to achieve this payoff is by partitioning into smaller groups.

A payoff vector $\bar{u}$ satisfies the *equal-treatment property* if $\bar{u}^{tq} = \bar{u}^{tq'}$ for all $q, q' \in \{1, \dots, n_t\}$ and for each $t = 1, \dots, T$.

Let $[n, (T, \Psi)]$ be a game and let β be a collection of subprofiles of n . The collection is a *balanced collection of subprofiles of n* if there are positive real numbers γ_s for $s \in \beta$ such that $\sum_{s \in \beta} \gamma_s s = n$. The numbers γ_s are called *balancing weights*. Given real number $\varepsilon \geq 0$, the game $[n; (T, \Psi)]$ is ε -*balanced* if for every balanced collection β of subprofiles of n it holds that

$$\Psi^*(n) \geq \sum_{s \in \beta} \gamma_s (\Psi(s) - \varepsilon \|s\|) \quad (7)$$

where the balancing weights for β are given by γ_s for $s \in \beta$. This definition extends that of Bondareva (1963) and Shapley (1967) to games with player types. Roughly, a game is (ε) balanced if allowing “part time” groups does not improve the total payoff (by more than ε per player). A game $[n; (T, \Psi)]$ is *totally balanced* if every subgame $[s; (T, \Psi)]$ is balanced.

The *balanced cover game* generated by a game $[n; (T, \Psi)]$ is a game $[n; (T, \Psi^b)]$ where

1. $\Psi^b(s) = \Psi(s)$ for all $s \neq n$ and
2. $\Psi^b(n) \geq \Psi(n)$ and $\Psi^b(n)$ is as small as possible consistent with the nonemptiness of the core of $[n; (T, \Psi^b)]$.

From the Bondareva-Shapley Theorem it follows that $\Psi^b(n) = \Psi^*(n)$ if and only if the game $[n; (T, \Psi)]$ is balanced (ε -balanced, with $\varepsilon = 0$).

For later convenience, the notion of the balanced cover of a pregame is introduced. Let (T, Ψ) be a pregame. For each profile s , define

$$\Psi^b(s) \stackrel{\text{def}}{=} \max_{\beta} \sum_{s \in \beta} \gamma_s \Psi(g), \quad (8)$$

where the maximum is taken over all balanced collections β of subprofiles of s with weights γ_g for $g \in \beta$. The pair (T, Ψ^b) is called the *balanced*

cover pregame of (T, Ψ) . Since a partition of a profile is a balanced collection it is immediately clear that $\Psi^b(s) \geq \Psi^*(s)$ for every profile s .

Premarkets

In this section, we introduce the concept of a premarket and re-state results from Shapley and Shubik (1969) in the context of pregames and premarkets.

Let $L + 1$ be a number of types of commodities and let $\{\hat{u}_t(y, \xi): t = 1, \dots, T\}$ denote a finite number of functions, called *utility functions*, of the form

$$\hat{u}_t(y, \xi) = u_t(y) + \xi,$$

where $y \in \mathbb{R}_+^L$ and $\xi \in \mathbb{R}$. (Such functions, in the literature of economics, are commonly called *quasi-linear*). Let $\{a^t \in \mathbb{R}_+^L: t = 1, \dots, T\}$ be interpreted as a set of *endowments*. We assume that $u_t(a^t) \geq 0$ for each t . For $t = 1, \dots, T$ we define $c^t \stackrel{\text{def}}{=} (u_t(\cdot), a^t)$ as a *participant type* and let $\mathbb{C} = \{c^t: t = 1, \dots, T\}$ be the set of participant types. Observe that from the data given by $\mathbb{C}$ we can construct a market by specifying a set of participants N and a function from N to $\mathbb{C}$ assigning endowments and utility functions – types – to each participant in N . A *premarket* is a pair $(T, \mathbb{C})$.

Let $(T, \mathbb{C})$ be a premarket and let $s = (s_1, \dots, s_T) \in \mathbb{Z}_+^T$. We interpret s as representing a group of economic participants with s_t participants having utility functions and endowments given by c^t for $t = 1, \dots, T$; for each t , that is, there are s_t participants in the group with type c^t . Observe that the data of a premarket gives us sufficient data to generate a pregame. In particular, given a profile $s = (s_1, \dots, s_T)$ listing numbers of participants of each of T types, define

$$W_{(s)} \stackrel{\text{def}}{=} \max \sum_t s_t u_t(y^t)$$

where the maximum is taken over the set $\{y^t \in \mathbb{R}_+^L: t = 1, \dots, T \text{ and } \sum_t s_t y^t = \sum_t a^t y^t\}$. Then the pair (T, W) is a *pregame generated by the premarket*.

The following Theorem is an extension to premarkets or a restatement of a result due to Shapley and Shubik (1969).

Theorem 1 Let $(T, \mathbb{C})$ be a premarket derived from economic data in which all utility functions are concave. Then the pregame generated by the premarket is totally balanced.

Direct Markets and Market-Game Equivalence

Shapley and Shubik (1969) introduced the notion of a direct market derived from a totally balanced game. In the direct market, each player is endowed with one unit of a commodity (himself) and all players in the economy have the same utility function. In interpretation, we might think of this as a labor market or as a market for productive factors, (as in (Owen 1975), for example) where each player owns one unit of a commodity. For games with player types as in this essay, we take the player types of the game as the commodity types of a market and assign all players in the market the same utility function, derived from the worth function of the game.

Let (T, Ψ) be a pregame and let $[n; (T, \Psi)]$ be a derived game. Let $N = \{(t, q): t = 1, \dots, T \text{ and } q = 1, \dots, n_t \text{ for each } t\}$ denote the set of players in the game where all participants $\{(t', q): q = 1, \dots, n_{t'}\}$ are of type t' for each $t' = 1, \dots, T$. To construct the direct market generated by a derived game $[n; (T, \Psi)]$, we take the commodity space as $\mathbb{R}_+^T$ and suppose that each participant in the market of type t is endowed with one unit of the t th commodity, and thus has endowment $1_t = (0, \dots, 0, 1, 0, \dots, 0) \in \mathbb{R}_+^T$ where “1” is in the t th position. The total endowment of the economy is then given by $\sum_t n_t 1_t = n$.

For any vectors $y \in \mathbb{R}_+^T$ define

$$u(y) \stackrel{\text{def}}{=} \max_{s \leq n} \sum_s \gamma_s \Psi(s), \quad (9)$$

the maximum running over all $\{\gamma_s \geq 0: s \in \mathbb{Z}_+^T, s \leq n\}$ satisfying

$$\sum_{s \leq n} \gamma_s s = y. \quad (10)$$

As noted by Shapley and Shubik (1969), but for our types case, it can be verified that the function u is concave and one-homogeneous. This does not depend on the balancedness of the game $[n; (T, \Psi)]$. Indeed, one may think of u as the

“balanced cover of $[n; (T, \Psi)]$ extended to R_+^T ”. Note also that u is superadditive, independent of whether the pregame (T, Ψ) is superadditive. We leave it to the interested reader to verify that if Ψ were not necessarily superadditive and Ψ^* is the superadditive cover of Ψ then it holds that $\max_{\sum s \leq n} \gamma_s \Psi(s) = \max_{\sum s \leq n} \gamma_s \Psi^*(s)$.

Taking the utility function u as the utility function of each player $(t, q) \in N$ where N is now interpreted as the set of participants in a market, we have generated a market, called the *direct market*, denoted by $[n, u; (T, \Psi)]$, from the game $[n; (T, \Psi)]$.

Again, the following extends a result of Shapley and Shubik (1969) to pregames.

Theorem 2 Let $[n, u; (T, \Psi)]$ denote the direct market generated by a game $[n; (T, \Psi)]$ and let $[n; (T, u)]$ denote the game derived from the direct market. Then, if $[n; (T, \Psi)]$ is a totally balanced game, it holds that $[n; (T, u)]$ and $[n; (T, \Psi)]$ are identical.

Remark 2 If the game $[n; (T, \Psi)]$ and every subgame $[s, (T, \Psi)]$ has a nonempty core – that is, if the game is ‘totally balanced’ then the game $[n; (T, u)]$ generated by the direct market is the initially given game $[n; (T, \Psi)]$. If however the game $[n; (T, \Psi)]$ is not totally balanced then $u(s) \geq \Psi(s)$ for all profiles $s \leq n$. But, whether or not $[n; (T, \Psi)]$ is totally balanced, the game $[n; (T, n)]$ is totally balanced and coincides with the totally balanced cover of $[n; (T, \Psi)]$.

Remark 3 Another approach to the equivalence of markets and games is taken by Garratt and Qin (1997), who define a class of direct lottery markets. While a player can participate in only one coalition, both ownership of coalitions and participation in coalitions is determined randomly. Each player is endowed with one unit of probability, his own participation. Players can trade their endowments at market prices. The core of the game is equivalent to the equilibrium of the direct market lottery.

Equivalence of Markets and Games with Many Players

The requirement of Shapley and Shubik (1969) that utility functions be concave is restrictive. It

rules out, for example situations such as economies with indivisible commodities. It also rules out club economies; for a given club structure of the set of players – in the simplest case, a partition of the total player set into groups where collective activities only occur within these groups – it may be that utility functions are concave over the set of alternatives available within each club, but utility functions need not be concave over all possible club structures. This rules out many examples; we provide a simple one below.

To obtain the result that with many players, games derived from pregames are market games, we need some further assumption on pregames. If there are many substitutes for each player, then the simple condition that *per capita payoffs are bounded* – that is, given a pregame (T, Ψ) , that there exists some constant K such that $\frac{\Psi(s)}{\|s\|} < K$ for all profiles s – suffices. If, however, there may be ‘scarce types’, that is, players of some type (s) become negligible in the population, then a stronger assumption of ‘small group effectiveness’ is required. We discuss these two conditions in the next section.

Small Group Effectiveness and Per Capita Boundedness

This section discusses conditions limiting gains to group size and their relationships. This definition was introduced in Wooders (1983), for NTU, as well as TU, games.

PCB A pregame (T, Ψ) satisfies *per capita boundedness* (PCB) if

$$PCB : \sup_{s \in Z_+^T} \frac{\Psi(s)}{\|s\|} \text{ is finite} \quad (11)$$

or equivalently,

$$\sup_{s \in Z_+^T} \frac{\Psi^*(s)}{\|s\|} \text{ is finite.}$$

It is known that under the apparently mild conditions of PCB and essential superadditivity, in general games with many players of each of a finite number of player types and a fixed distribution of player types have nonempty approximate cores; Wooders (1977, 1983). (Forms of these

assumptions were subsequently also used in Shubik and Wooders (1982c, 1983b); Kaneko and Wooders (1986); and Wooders (1992b, 1994) among others.) Moreover, under the same conditions, approximate cores have the property that most players of the same type are treated approximately equally (Wooders 1977, 2008a; see also Shubik and Wooders (1982c)). These results, however, either require some assumption ruling out ‘scarce types’ of players, for example, situations where there are only a few players of some particular type and these players can have great effects on total feasible payoffs. Following are two examples. The first illustrates that PCB does not control limiting properties of the per capita payoff function when some player types are scarce.

Example 3 (Wooders 2008a) Let $T = 2$ and let (T, Ψ) be the pregame given by

$$\Psi(s_1, s_2) = \begin{cases} s_1 + s_2 & \text{when } s_1 > 0 \\ 0 & \text{otherwise.} \end{cases}$$

The function Ψ obviously satisfies PCB. But there is a problem in defining $\lim \Psi(s_1, s_2)/s_1 + s_2$ as $s_1 + s_2$ tends to infinity, since the limit depends on how it is approached. Consider the sequence (s_1^v, s_2^v) where $(s_1^v, s_2^v) = (0, v)$; then $\lim \Psi(s_1^v, s_2^v)/s_1^v, s_2^v = 0$. Now suppose in contrast that $(s_1^v, s_2^v) = (1, v)$; then $\lim \Psi(s_1^v, s_2^v)/s_1^v, s_2^v = 1$. This illustrates why, to obtain the result that games with many players are market games either it must be required that there are no scarce types or some assumption limiting the effects of scarce types must be made. We return to this example in the next section.

The next example illustrates that, with only PCB, uniform approximate cores of games with many players derived from pregames may be empty.

Example 4 (Wooders 2008a) Consider a pregame (T, Ψ) where $T = \{1, 2\}$ and Ψ is the superadditive cover of the function Ψ' defined by:

$$\Psi'(s) \stackrel{\text{def}}{=} \begin{cases} |s| & \text{if } s_1 = 2, \\ 0 & \text{otherwise.} \end{cases}$$

Thus, if a profiles $= (s_1, s_2)$ has $s_1 = 2$ then the worth of the profile according to Ψ' is equal to the total number of players it represents, $s_1 + s_2$, while

all other profiles s have worth of zero. In the superadditive cover game the worth of a profile s is 0 if $s_1 < 2$ and otherwise is equal to s_2 plus the largest even number less than or equal to s_1 .

Now consider a sequence of profiles $(s^v)_v$ where $s_1^v = 3$ and $s_2^v = v$ for all v . Given $\varepsilon > 0$, for all sufficiently large player sets the uniform ε -core is empty. Take, for example, $\varepsilon = 1/4$. If the uniform ε -core were nonempty, it would have to contain an equal-treatment payoff vector. (It is well known and easily demonstrated that the uniform ε -core of a TU game is nonempty if and only if it contains an equal treatment payoff vector. This follows from the fact that the uniform ε -core is a convex set.) For the purpose of demonstrating a contradiction, suppose that $u^v = (u_1^v, u_2^v)$ represents an equal treatment payoff vector in the uniform ε -core of $[s^v; (T, \Psi)]$. The following inequalities must hold:

$$\begin{aligned} 3u_1^v + vu_2^v &\leq v + 2, \\ 2u_1^v + vu_2^v &\geq v + 2, \text{ and} \\ u_1^v &\geq \frac{3}{4}. \end{aligned}$$

which is impossible. A payoff vector which assigns each player zero is, however, in the weak ε -core for any $\varepsilon > \frac{1}{v+3}$. But it is not very appealing, in situations such as this, to ignore a relatively small group of players (in this case, the players of type 1) who can have a large effect on per capita payoffs. This leads us to the next concept.

To treat the scarce types problem, Wooders (1992a, b, 1993) introduced the condition of small group effectiveness (SGE). SGE is appealing technically since it resolves the scarce types problem. It is also economically intuitive and appealing; the condition defines a class of economies that, when there are many players, generate competitive markets. Informally, SGE dictates that *almost all* gains to collective activities can be realized by relatively small groups of players. Thus, SGE is exactly the sort of assumption required to ensure that multiple, relatively small coalitions, firms, jurisdictions, or clubs, for example, are optimal or near-optimal in large economies.

A pregame (T, Ψ) satisfies *small group effectiveness*, SGE, if:

For each real number $\varepsilon > 0$,
 there is an integer $\eta_0(\varepsilon)$
 $\times$ such that for each profiles,
 SGE : for some partition $\{s^k\}$ of s with
 $\|s^k\| \leq \eta_0(\varepsilon) \times$ for each subprofile s^k , it holds that
 $\Psi^*(s) - \sum_k \Psi(s^k) \leq \varepsilon \|s\|$;

(12)

given $\varepsilon > 0$ there is a group size $\eta_0(\varepsilon)$ such that the loss from restricting collective activities within groups to groups containing fewer than $\eta_0(\varepsilon)$ members is at most ε per capita (Wooders 1992a). (Exactly the same definition applies to situations with a compact metric space of player types, c.f. Wooders (1988, 1992a).)

SGE also has the desirable feature that if there are no ‘scarce types’ – types of players that appear in vanishingly small proportions – then SGE and PCB are equivalent.

Theorem 3 (Wooders 1994) *With ‘thickness,’ SGE = PCB* (1) Let (T, Ψ) be a pregame satisfying SGE. Then the pregame satisfies PCB.

(2) Let (T, Ψ) be a pregame satisfying PCB. Then given any positive real number ρ , construct a new pregame (T, Ψ_ρ) where the domain of Ψ_ρ is restricted to profiles s where, for each $t = 1, \dots, T$, either $\frac{s_t}{\|s\|} > \rho$ or $s_t = 0$. Then (T, Ψ_ρ) satisfies SGE on its domain.

It can also be shown that small groups are effective for the attainment of nearly all feasible outcomes, as in the above definition, if and only if small groups are effective for improvement – any payoff vector that can be significantly improved upon can be improved upon by a small group (see Proposition 3.8 in Wooders 1992b).

Remark 4 Under a stronger condition of *strict* small group effectiveness, which dictates that $\eta(\varepsilon)$ in the definition of small group effectiveness can be chosen independently of ε , stronger results can be obtained than those presented in this section and the next. We refer to Winter and Wooders (1990) for a treatment of this case.

Remark 5 (On the importance of taking into account scarce types) Recall the quotation from von Neumann and Morgenstern and the discussion following the quotation. The assumption of

per capita boundedness has significant consequences but is quite innocuous – ruling out the possibility of average utilities becoming infinite as economies grow large does not seem restrictive. But with only per capita boundedness, even the formation of small coalitions can have significant impacts on aggregate outcomes. With small group effectiveness, however, there is no problem of either large or small coalitions acting together – large coalitions cannot do significantly better than relatively small coalitions.

Roughly, the property of large games we next introduce is that relatively small groups of players make only “asymptotic negligible” contributions to per-capita payoffs of large groups. A pregame (Ω, Ψ) satisfies *asymptotic negligibility* if, for any sequence of profiles $\{f^v\}$ where

$$\begin{aligned} \|f^v\| &\rightarrow \infty \text{ as } v \rightarrow \infty, \\ \sigma(f^v) &= \sigma(f^{v'}) \text{ for all } v \text{ and } v' \text{ and} \\ \lim_{v \rightarrow \infty} \frac{\Psi^*(f^v)}{\|f^v\|} &\text{ exists,} \end{aligned} \quad (13)$$

then for any sequence of profiles $\{\ell^v\}$ with

$$\lim_{v \rightarrow \infty} \frac{\|\ell^v\|}{\|f^v\|} = 0, \quad (14)$$

it holds that

$$\begin{aligned} \lim_{v \rightarrow \infty} \frac{\Psi^*(\|f^v + \ell^v\|)}{\|f^v + \ell^v\|} &\text{ exists, and} \\ \lim_{v \rightarrow \infty} \frac{\Psi^*(\|f^v + \ell^v\|)}{\|f^v + \ell^v\|} &= \lim_{v \rightarrow \infty} \frac{\Psi^*(f^v)}{\|f^v\|}. \end{aligned} \quad (15)$$

Theorem 4 (Wooders 1992b, 2008b) A pregame (T, Ψ) satisfies SGE if and only if it satisfies PCB and asymptotic negligibility. Intuitively, asymptotic negligibility ensures that vanishingly small percentages of players have vanishingly small effects on aggregate per-capita worths. It may seem paradoxical that SGE, which highlights the importance of relatively small groups, is equivalent to asymptotic negligibility. To gain some intuition, however, think of a marriage model where only two-person marriages are allowed. Obviously two-person groups are (strictly) effective, but also, in large player sets, no two persons can have a substantial affect on aggregate per-capita payoffs.

Remark 6 Without some assumptions ensuring essential superadditivity, at least as incorporated into our definition of feasibility, nonemptiness of approximate cores of large games cannot be expected; superadditivity assumptions (or the close relative, essential superadditivity) are heavily relied upon in all papers on large games cited. In the context of economies, superadditivity is a sort of monotonicity of preferences or production functions assumption, that is, superadditivity of Ψ implies that for all $s, s' \in Z_+^T$, it holds that $\Psi(s + s') \geq \Psi(s) + \Psi(s')$. Our assumption of small group effectiveness, SGE, admits non-monotonicities. For example, suppose that ‘two is company, three or more is a crowd,’ by supposing there is only one commodity and by setting $\Psi(2) = 2$, $\Psi(n) = 0$ for $n \neq 2$. The reader can verify, however, that this example satisfies small group effectiveness since $\Psi^*(n) = n$ if n is even and $\Psi^*(n) = n - 1$ otherwise. Within the context of pregames, requiring the superadditive cover payoff to be approximately realizable by partitions of the total player set into relatively small groups is the weakest form of superadditivity required for the equivalence of games with many players and concave markets.

Derivation of Markets from Pregames Satisfying SGE

With SGE and PCB in hand, we can now derive a premarket from a pregame and relate these concepts.

To construct a limiting direct premarket from a pregame, we first define an appropriate utility function. Let (T, Ψ) be a pregame satisfying SGE. For each vector x in R_+^T define

$$U(x) \stackrel{\text{def}}{=} \|x\| \lim_{v \rightarrow \infty} \frac{\Psi^*(f^v)}{\|f^v\|} \quad (16)$$

where the sequence $\{f^v\}$ satisfies

$$\lim_{v \rightarrow \infty} \frac{f^v}{\|f^v\|} = \frac{x}{\|x\|} \quad \text{and} \quad \|f^v\| \rightarrow \infty \quad (17)$$

Theorem 5 (Wooders 1988, 1994) Assume the pregame (T, Ψ) satisfies small group effectiveness.

Then for any $x \in R_+^T$ the limit (16) exists. Moreover, $U(\cdot)$ is well-defined, concave and 1-homogeneous and the convergence is uniform in the sense that, given $\varepsilon > 0$ there is an integer η such that for all profiles s with $\|s\| \leq \eta$ it holds that

$$\left| U\left(\frac{s}{\|s\|}\right) - \frac{\Psi^*(s)}{\|s\|} \right| \leq \varepsilon.$$

From Wooders (1994) (Theorem 4), if arbitrarily small percentages of players of any type that appears in games generated by the pregame are ruled out, then the above result holds under per capita boundedness (Wooders 1994) (Theorem 6). As noted in the introduction to this paper, for the TU case, the concavity of the limiting utility function, for the model of Wooders (1983) was first noted by Aumann (1987). The concavity is shown to hold with a compact metric space of player types in Wooders (1988) and is simplified to the finite types case in Wooders (1994).

Theorem 5 follows from the facts that the function U is superadditive and 1-homogeneous on its domain. Since U is concave, it is continuous on the interior of its domain; this follows from PCB. Small group effectiveness ensures that the function U is continuous on its entire domain (Wooders 1994) (Lemma 2).

Theorem 6 (Wooders 1994) Let (T, Ψ) be a pregame satisfying small group effectiveness and let (T, U) denote the derived direct market pregame. Then (T, U) is a totally balanced market game. Moreover, U is one-homogeneous, that is, $U(\lambda x) = \lambda U(x)$ for any non-negative real number λ .

In interpretation, T denotes a number of types of players/commodities and U denotes a utility function on R_+^T . Observe that when U is restricted to profiles (in Z_+^T), the pair (T, U) is a pregame with the property that every game $[n; (T, U)]$ has a nonempty core; thus, we will call (T, U) the *premarket generated by the pregame* (T, Ψ) . That every game derived from (T, U) has a nonempty core is a consequence of the Shapley and Shubik (1969) result that market games derived from markets with concave utility functions are totally balanced.

It is interesting to note that, as discussed in Wooders (Sect. 6 in Wooders 1994), if we restrict the number of commodities to equal the number of player types, then the utility function U is *uniquely* determined. (If one allowed more commodities then one would effectively have ‘redundant assets’.) In contrast, for games and markets of fixed, finite size, as demonstrated in Shapley and Shubik (1975), even if we restrict the number of commodities to equal the number of player types, given any nonempty, compact, convex subset of payoff vectors in the core, it is possible to construct utility functions so that this subset coincides with the set of competitive payoffs. Thus, in the Shapley and Shubik approach, equivalence of the core and the set of price-taking competitive outcomes for the direct market is only an artifact of the method used there of constructing utility functions from the data of a game and is quite distinct from the equivalence of the core and the set of competitive payoff vectors as it is usually understood (that is, in the sense of Debreu and Scarf (1963) and Aumann (1964). See also Kalai and Zemel (1982a, b) which characterize the core in multi-commodity flow games.

Cores and Approximate Cores

The concept of the core clearly was important in the work of Shapley and Shubik (1966, 1969, 1975) and is also important for the equivalence of games with many players and market games. Thus, we discuss the related results of non-emptiness of approximate cores and convergence of approximate cores to the core of the ‘limit’ – the game where all players have utility functions derived from a pregame and large numbers of players. First, some terminology is required. A vector p is a *subgradient* at x of the concave function U if $U(y) - U(x) \leq p \cdot (y - x)$ for all y . One might think of a subgradient as a bounding hyperplane. To avoid any confusion it might be helpful to note that, as Mas-Colell (1985) remarks: “Strictly speaking, one should use the term *subgradient* for convex functions and *supergradient* for concave. But this is cumbersome”, (pp. 29–30 in Mas-Colell 1985).

For ease of notation, equal-treatment payoff vectors for a game $[n; (T, \Psi)]$ will typically be represented as vectors in $\mathbb{R}^T$. An *equal-treatment payoff vector*, or simply a *payoff vector* when the meaning is clear, is a point $\bar{x}$ in $\mathbb{R}^T$. The t^{th} component of $\bar{x}$, $\bar{x}_t$, is interpreted as the payoff to each player of type t . The feasibility of an equal-treatment payoff vector $\bar{x} \in \mathbb{R}^T$ for the game $[n; (T, \Psi)]$ can be expressed as:

$$\Psi^*(n) \geq \bar{x} \cdot n.$$

Let $[n; (T, \Psi)]$ be a game determined by a pregame (T, Ψ) , let ε be a non-negative real number, and let $\bar{x} \in \mathbb{R}^T$ be a (equal-treatment) payoff vector. Then $\bar{x}$ is in the *equal-treatment ε -core* of $[n; (T, \Psi)]$ or simply “in the ε -core” when the meaning is clear, if x is feasible for $[n; (T, \Psi)]$ and

$$\Psi(s) \leq \bar{x} \cdot s + \varepsilon \|s\| \text{ for all subprofiles } s \text{ of } n.$$

Thus, the equal-treatment ε -core is the set

$$C(n; \varepsilon) \stackrel{\text{def}}{=} \{ \bar{x} \in \mathbb{R}_+^T : \Psi^*(n) \geq \bar{x} \cdot n \text{ and } \Psi(s) \leq \bar{x} \cdot s + \varepsilon \|s\| \text{ for all subprofiles } s \text{ of } n \}. \quad (18)$$

It is well known that the ε -core of a game with transferable utility is nonempty if and only if the equal-treatment ε -core is nonempty.

Continuing with the notation above, for any $s \in R_+^T$, let $\Pi(s)$ denote the set of subgradients to the function U at the point s ;

$$\Pi(s) \stackrel{\text{def}}{=} \{ \pi \in \mathbb{R}^T : \pi \cdot s \text{ and } \pi \cdot s' \geq U(s') \text{ for all } s' \in R_+^T \}. \quad (19)$$

The elements in $\Pi(s)$ can be interpreted as equal-treatment core payoffs to a limiting game with the mass of players of type t given by s_t . The core payoff to a player is simply the value of the one unit of a commodity (himself and all his attributes, including endowments of resources) that he owns in the direct market generated by a game. Thus $\Pi(\cdot)$ is called the limiting core correspondence for the pregame (T, Ψ) . Of course $\Pi(\cdot)$ is also the limiting core correspondence for the pregame (T, U) .

Let $\hat{\Pi}(n) \subset \mathbb{R}^T$ denote *equal-treatment core of the market game* $[n; (T, u)]$.

$$\hat{\Pi}(n) \stackrel{\text{def}}{=} \left\{ \pi \in \mathbb{R}^T : \pi \cdot n = u(n) \right. \\ \left. \text{and } \pi \cdot s \geq u(s) \text{ for all } s \in \mathbb{Z}_+^T, s \leq n \right\}. \quad (20)$$

Given any player profile n and derived games $[n; (T, \Psi)]$ and $[n; (T, U)]$ it is interesting to observe the distinction between the equal-treatment core of the game $[n; (T, U)]$, denoted by $\Pi(n)$, defined by (20), and the set $U(n)$ (that is, $\Pi(x)$ with $x = n$). The definitions of $\Pi(n)$ and $\hat{\Pi}(n)$ are the same except that the qualification “ $s \leq n$ ” in the definition of $\hat{\Pi}(n)$ does not appear in the definition of $\Pi(n)$. Since $U(n)$ is the *limiting core* correspondence, it takes into account arbitrarily large coalitions. For this reason, for any $x \in \Pi(n)$ and $\hat{x} \in \hat{\Pi}(n)$ it holds that $\bar{x} \cdot n \geq \hat{x} \cdot n$. A simple example may be informative.

Example 5 Let (T, Ψ) be a pregame where $T = 1$ and $\Psi(n) = n - \frac{1}{n}$ for each $n \in \mathbb{Z}_+$, and let $[n; (T, \Psi)]$ be a derived game. Then $\Pi(n) = \{1\}$ while $\hat{\Pi}(n) = \left\{ \left(1 - \frac{1}{n^2} \right) \right\}$.

The following Theorem extends a result due to Shapley and Shubik (1975) stated for games derived from pregames.

Theorem 7 (Shapley and Shubik 1975) Let $[n; (T, \Psi)]$ be a game derived from a pregame and let $[n, u; (T, \Psi)]$ be the direct market generated by $[n; (T, \Psi)]$. Then the equal-treatment core $\Pi(n)$ of the game $[n; (T, u)]$ is nonempty and coincides with the set of competitive price vectors for the direct market $[n, u; (T, \Psi)]$.

Remark 7 Let (T, Ψ) be a pregame satisfying PCB. In the development of the theory of large games as models of competitive economies, the following function on the space of profiles plays an important role:

$$\lim_{r \rightarrow \infty} \frac{\Psi^*(rf)}{r};$$

see, for example, Wooders (1977) and Shubik and Wooders (1982c). For the purposes of

comparison, we introduce another definition of a limiting utility. Function For each vector x in $\mathbb{R}_+^T$ with rational components let $r(x)$ be the smallest integer such that $r(x)x$ is a vector of integers. Therefore, for each rational vector x , we can define

$$\hat{U}(x) \stackrel{\text{def}}{=} \lim_{v \rightarrow \infty} \frac{\Psi^*(vr(x)x)}{vr(x)}.$$

Since Ψ^* is superadditive and satisfies per capita boundedness, the above limit exists and $\hat{U}(\cdot)$ is well-defined. Also, $\hat{U}(x)$ has a continuous extension to any closed subset strictly in the interior of $\mathbb{R}_+^T$. The function $\hat{U}(x)$, however, may be discontinuous at the boundaries of $\mathbb{R}_+^T$. For example, suppose that $T = 2$ and

$$\Psi^*(k, n) = \begin{cases} k + 1 & \text{when } k > 0 \\ 0 & \text{otherwise.} \end{cases}$$

The function Ψ^* obviously satisfies PCB but does not satisfy SGE. To see the continuity problem, consider the sequences $\{x^v\}$ and $\{y^v\}$ of vectors in $\mathbb{R}_+^T$ where $x^v = (\frac{1}{v}, \frac{v-1}{v})$ and $y^v = (0, v)$. Then $\lim_{v \rightarrow \infty} x^v = \lim_{v \rightarrow \infty} y^v = (0, 1)$ but $\lim_{v \rightarrow \infty} \hat{U}(x^v) = 1$ while $\lim_{v \rightarrow \infty} \hat{U}(y^v) = 0$. SGE is precisely the condition required to avoid this sort of discontinuity, ensuring that the function U is continuous on the boundaries of $\mathbb{R}_+^T$.

Before turning to the next section, let us provide some additional interpretation for $\hat{\Pi}(n)$. Suppose a game $[n; (T, \Psi)]$ is one generated by an economy, as in Shapley and Shubik (1966) or Owen (1975), for example. Players of different types may have different endowments of private goods. An element π in $\hat{\Pi}(n)$ is an equal-treatment payoff vector in the core of the balanced cover game generated by $[n; (T, \Psi)]$ and can be interpreted as listing prices for player types where π_t is the price of a player of type t ; this price is a price for the player himself, *including* his endowment of private goods.

Nonemptiness and Convergence of Approximate Cores of Large Games

The next Proposition is an immediate consequence of the convergence of games to markets

shown in Wooders (1992b, 1994) and can also be obtained as a consequence of Theorem 5 above.

Proposition 2 (*Nonemptiness of approximate cores*) Let (T, Ψ) be a pregame satisfying SGE. Let ε be a positive real number. Then there is an integer $\eta_1(\varepsilon)$ such that any game $[n; (T, \Psi)]$ with $\|n\| \geq \eta_1(\varepsilon)$ has a nonempty uniform ε -core.

(Note that no assumption of superadditivity is required but only because our definition of feasibility is equivalent to feasibility for superadditive covers.)

The following result was stated in Wooders (1992b). For more recent results see Wooders (2008a).

Theorem 8 ((Wooders 1992b) *Uniform closeness of (equal-treatment) approximate cores to the core of the limit game*) Let (T, Ψ) be a pregame satisfying SGE and let $\Pi(\cdot)$ be as defined above. Let $\delta > 0$ and $\rho > 0$ be positive real numbers. Then there is a real number ε^* with $0 < \varepsilon^*$ and an integer $\eta_0(\delta, \rho, \varepsilon^*)$ with the following property: for each positive $\varepsilon \in (0, \varepsilon^*)$ and each game $[f; (T, \Psi)]$ with $\|f\| > \eta_0(\delta, \rho, \varepsilon^*)$ and $f_t/\|f\| \geq \rho$ for each $t = 1, \dots, T$, if $C(f; \varepsilon)$ is nonempty then both

$$\text{dist}[C(f; \varepsilon), \Pi(f)] < \delta \text{ and } \text{dist}[C(f; \varepsilon), \hat{\Pi}(f)] < \delta,$$

where ‘dist’ is the Hausdorff distance with respect to the sum norm on $\mathbb{R}^T$.

Note that this result applies to games derived from diverse economies, including economies with indivisibilities, nonmonotonicities, local public goods, clubs, and so on.

Theorem 8 motivates the question of whether approximate cores of games derived from pregames satisfying small group effectiveness treat players most of the same type nearly equally. The following result, from Wooders (1977, 1992b, 2007) answers this question.

Theorem 9 Let (T, Ψ) be a pregame satisfying SGE. Then given any real numbers $\gamma > 0$ and $\lambda > 0$ there is a positive real number ε^* and an integer ρ such that for each $\varepsilon \in [0, \varepsilon^*)$ and for every profile $n \in \mathbb{Z}_+^T$ with $\|n\|_1 > \rho$, if $x \in \mathbb{R}^N$ is in the uniform ε -core of the game $[n, \Psi]$ with player set

$$N = \{(t, q) : t = 1, \dots, T, \text{ and for each } t, q = 1, \dots, n_t\}$$

then, for each $t \in \{1, \dots, T\}$ with $\frac{n_t}{\|n\|_1} \geq \frac{\lambda}{2}$ it holds that

$$|\{t, q\} : |x^{tq} - z_t| > \gamma\}| < \lambda n_t,$$

where, for each $t = 1, \dots, T$,

$$z_t = \frac{1}{n} \sum_{q=1}^{n_t} x^{tq},$$

the average payoff received by players of type t .

Shapley Values of Games with Many Players

Let (N, v) be a game. The *Shapley value* of a superadditive game is the payoff vector whose i th component is given by

$$SH(v, i) = \frac{1}{|N|} \sum_{j=0}^{|N|-1} \frac{1}{\binom{|N|-1}{j}} \sum_{\substack{S \subset N \setminus \{i\} \\ |S|=j}} [v(S \cup \{i\}) - v(S)].$$

To state the next Theorem, we require one additional definition. Let (T, Ψ) be a pregame. The pregame satisfies *boundedness of marginal contributions* (BMC) if there is a constant M such that

$$|\Psi(s + 1_t) - \Psi(s)| \leq M$$

for all vectors $1_t = (0, \dots, 0, 1_t^{\text{th place}}, 0, \dots, 0)$ for each $t = 1, \dots, T$. Informally, this condition bounds marginal contributions while SGE bounds average contributions. That BMC implies SGE is shown in Wooders (1992b). The following result restricts the main Theorem of Wooders and Zame (1987) to the case of a finite number of types of players.

Theorem 10 (Wooders and Zame 1987) Let (T, Ψ) be a superadditive pregame satisfying boundedness of marginal contributions. For each $\varepsilon > 0$ there is a number $\delta(\varepsilon) > 0$ and an integer $\mu(\varepsilon)$ with the following property:

If $[n, (T, \Psi)]$ is a game derived from the pregame, for which $n_t > \mu(\varepsilon)$ for each t , then the Shapley value of the game is in the (weak) ε -core.

Similar results hold within the context of private goods exchange economies (cf., Shapley (1964), Shapley and Shubik (1969), Champsaur (1975), Mas-Colell (1977), Cheng (1981) and others). Some of these results are for economies without money but all treat private goods exchange economies with divisible goods and concave, monotone utility functions. Moreover, they all treat either replicated sequences of economies or convergent sequences of economies. That games satisfying SGE are asymptotically equivalent to balanced market games clarifies the contribution of the above result. In the context of the prior results developed in this paper, the major shortcoming of the Theorem is that it requires BMC. This author conjectures that the above result, or a close analogue, could be obtained with the milder condition of SGE, but this has not been demonstrated.

Economies with Clubs

By a club economy we mean an economy where participants in the economy form groups – called clubs – for the purposes of collective consumption and/or production collectively with the group members. The groups may possibly overlap. A club structure of the participants in the economy is a covering of the set of players by clubs. Providing utility functions are quasi-linear, such an economy generates a game of the sort discussed in this essay. The worth of a group of players is the maximum total worth that the group can achieve by forming clubs. The most general model of clubs in the literature at this point is Allouch and Wooders (2008). Yet, if one were to assume that utility functions were all quasi-linear and the set of possible types of participants were finite. The results of this paper would apply.

In the simplest case, the utility of an individual depends on the club profile (the numbers of participants of each type) in his club. The total worth of a group of players is the maximum that it can achieve by splitting into clubs. The results presented in this section immediately apply. When there are many participants, club economies can be represented as markets and the competitive payoff vectors for the market are approximated by equal-treatment payoff vectors in approximate cores. Approximate cores converge to equal treatment and competitive equilibrium payoffs. A more general model making these points is treated in Shubik and Wooders (1982a). For recent reviews of the literature, see Conley and Smith (2005) and Kovalenkov and Wooders (2005). (Other approaches to economies with clubs/local public goods include Casella and Feinstein (2002), Demange (1994), Haimanko, O., M. Le Breton and S. Weber (2004), and Konishi, Le Breton and Weber (1998). Recent research has treated clubs as networks.)

Coalition production economies may also be viewed as club economies. We refer the reader to Böhm (1974), Sondermann (1974), Shubik and Wooders (1983b), and for a more recent treatment and further references, Sun, Trockel and Yang (2008).

Let us conclude this section with some historical notes. Club economies came to the attention of the economics profession with the publication of Buchanan (Buchanan 1965). The author pointed out that people care about the numbers of other people with whom they share facilities such as swimming pool clubs. Thus, there may be congestion, leading people to form multiple clubs. Interestingly, much of the recent literature on club economies with many participants and their competitive properties has roots in an older paper, Tiebout (1956). Tiebout conjectured that if public goods are ‘local’ – that is, subject to exclusion and possibly congestion – then large economies are ‘market-like’. A first paper treating club economies with many participants was Pauly (1970), who showed that, when all players have the same preferred club size, then the core of economy is nonempty if and only if all participants in the economy can be partitioned into groups of the

preferred size. Wooders (1978) modeled a club economy as one with local public goods and demonstrated that, when individuals within a club (jurisdiction) are required to pay the same share of the costs of public good provision, then outcomes in the core permit heterogeneous clubs if and only if all types of participants in the same club have the same demands for local public goods and for congestion. Since these early results, the literature on clubs has grown substantially.

With a Continuum of Players

Since Aumann (1964) much work has been done on economies with a continuum of players. It is natural to question whether the asymptotic equivalence of markets and games reported in this article holds in a continuum setting. Some such results have been obtained.

First, let $N = [0, 1]$ be the 0,1 interval with Lebesgue measure and suppose there is a partition of N into a finite set of subsets $N_1, \dots, N_T$ where, in interpretation, a point in N_t represents a player of type t . Let Ψ be given. Observe that Ψ determines a payoff for any finite group of players, depending on the numbers of players of each type. If we can aggregate partitions of the total player set into finite coalitions then we have defined a game with a continuum of players and finite coalitions.

For a partition of the continuum into finite groups to ‘make sense’ economically, it must preserve the relative scarcities given by the measure. This was done in Kaneko and Wooders (1986). To illustrate their idea of measurement consistent partitions of the continuum into finite groups, think of a census form that requires each three-person household to label the players in the household, #1, #2, or #3. When checking the consistency of its figures, the census taker would expect the numbers of people labeled #1 in three-person households to equal the numbers labeled #2 and #3. For consistency, the census taker may also check that the number of first persons in three-person households in a particular region is equal to the number of second persons and third persons

in three person households in that region. It is simple arithmetic. This consistency should also hold for k -person households for any k . Measurement consistency is the same idea with the work “number” replaced by “proportion” or “measure”.

One can immediately apply results reported above to the special case of TU games of Kaneko-Wooders (1986) and conclude that games satisfying small group effectiveness and with a continuum of players have nonempty cores and that the payoff function for the game is one-homogeneous. (We note that there have been a number of papers investigating cores of games with a continuum of players that have come to the conclusion that non-emptiness of exact cores does not hold, even with balancedness assumptions, cf., Weber (1979, 1981)). The results of Wooders (1994), show that the continuum economy must be representable by one where all players have the same concave, continuous one-homogeneous utility functions. Market games with a continuum of players and a finite set of types are also investigated in Azriel and Lehrer (2007), who confirm these conclusions.)

Other Related Concepts and Results

In an unpublished 1972 paper due to Edward Zajac (1972), which has motivated a large amount of literature on ‘subsidy-free pricing’, cost sharing, and related concepts, the author writes:

“A fundamental idea of equity in pricing is that ‘no consumer group should pay higher prices than it would pay by itself. . .’. If a particular group is paying a higher price than it would pay if it were severed from the total consumer population, the group feels that it is subsidizing the total population and demands a price reduction”.

The “dual” of the cost allocation problem is the problem of surplus sharing and subsidy-free pricing. (See, for example Moulin (1988, 1992) for excellent discussions of these two problems.) Tauman (1987) provides a excellent survey. Some recent works treating cost allocation and subsidy free-pricing include Moulin (1988, 1992). See also the recent notion of “Walras’ core” in Qin, Shapley and Shimomura (2006).

Another related area of research has been into whether games with many players satisfy some notion of the Law of Demand of consumer theory (or the Law of Supply of producer theory). Since games with many players resemble market games, which have the property that an increase in the endowment of a commodity leads to a decrease in its price, such a result should be expected. Indeed, for games with many players, a Law of Scarcity holds – if the numbers of players of a particular type is increased, then core payoffs to players of that type do not increase and may decrease. (This result was observed by Scotchmer and Wooders (1988)). See Kovalenkov and Wooders (2005, 2006) for the most recent version of such results and a discussion of the literature. Laws of scarcity in economies with clubs are examined in Cartwright, Conley and Wooders (2006).

Some Remarks on Markets and More General Classes of Economies

Forms of the equivalence of outcomes of economies where individuals have concave utility functions but not necessarily linear in money. These include Billera (1974), Billera and Bixby (1974) and Mas-Colell (1975). A natural question is whether the results reported in this paper can extend to nontransferable utility games and economies where individuals have utility functions that are not necessarily linear in money. So far the results obtained are not entirely satisfactory. Nonemptiness of approximate cores of games with many players, however, holds in substantial generality; see Kovalenkov and Wooders (2003) and Wooders (2008b).

Conclusions and Future Directions

The results of Shapley and Shubik (1969), showing equivalence of structures, rather than equivalence of outcomes of solution concepts in a fixed structure (as in Aumann 1964, for example) are remarkable. So far, this line of research has been relatively little explored. The results for games with many players have also not been fully

explored, except for in the context of games, such as those derived from economies with clubs, and with utility functions that are linear in money.

Per capita boundedness seems to be about the mildest condition that one can impose on an economic structure and still have scarcity of per capita resources in economies with many participants. In economies with quasi-linear utilities (and here, I mean economies in a general sense, as in the glossary) satisfying per capita boundedness and where there are many substitutes for each type of participant, then as the number of participants grows, these economies resemble or (as if they) *are* market economies where individuals have continuous, and monotonic increasing utility functions. Large groups cannot influence outcomes away from outcomes in the core (and outcomes of free competition) since large groups are not significantly more effective than many small groups (from the equivalence, when each player has many close substitutes, between per capita boundedness and small group effectiveness).

But if there are not many substitutes for each participant, then, as we have seen, per capita boundedness allows small groups of participants to have large effects and free competition need not prevail (cores may be empty and price-taking equilibrium may not exist). The condition required to ensure free competition in economies with many participants, without assumptions of “thickness”, is precisely small group effectiveness.

But the most complete results relating markets and games, outlined in this paper, deal with economies in which all participants have utility functions that are linear in money and in games with side payments, where the worth of a group can be divided in any way among the members of the group without any loss of total utility or worth. Nonemptiness of approximate cores of large games without side payments has been demonstrated; see Wooders (1983, 2008b) and Kovalenkov and Wooders (2003). Moreover, it has been shown that when side payments are limited then approximate cores of games without side payments treat similar players similarly (Kovalenkov and Wooders 2001).

Results for *specific economic structures*, relating cores to price taking equilibrium treat can treat situations that are, in some respects, more general. A substantial body of literature shows that certain classes of club economies have nonempty cores and also investigates price-taking equilibrium in these situations. Fundamental results are provided by Gale and Shapley (1962), Shapley and Shubik (1972), and Crawford and Kelso (1982) and many more recent papers. We refer the reader to Roth and Sotomayor (1990) and to Two-Sided Matching Models, by Ömer and Sotomayor in this encyclopedia. A special feature of the models of these papers is that there are two sorts of players or two sides to the market; examples are (1) men and women, (2) workers and firms, (3) interns and hospitals and so on.

Going beyond two-sided markets to clubs in general, however, one observes that the positive results on nonemptiness of cores and existence of price-taking equilibria only holds under restrictive conditions. A number of recent contributions however, provide specific economic models for which, when there are many participants in the economy, as in exchange economies it holds that price-taking equilibrium exists, cores are non-empty, and the set of outcomes of price-taking equilibrium are equivalent to the core (see, for example, Allouch and Wooders 2008; Allouch et al. 2008; Ellickson et al. 1999; Wooders 1989, 1997).

Bibliography

- Allouch N, Wooders M (2008) Price taking equilibrium in economies with multiple memberships in clubs and unbounded club sizes. *J Econ Theor* 140:246–278
- Allouch N, Conley JP, Wooders M (2008) Anonymous price taking equilibrium in Tiebout economies with a continuum of agents: existence and characterization. *J Math Econ*. <https://doi.org/10.1016/j.jmateco.2008.06.003>
- Aumann RJ (1964) Markets with a continuum of traders. *Econometrica* 32:39–50
- Aumann RJ (1987) Game theory. In: Eatwell J, Milgate M, Newman P (eds) *The new Palgrave: a dictionary of economics*. Palgrave MacMillan, Basingstoke
- Aumann RJ, Dreze J (1974) Cooperative games with coalition structures. *Int J Game Theory* 3:217–237
- Aumann RJ, Shapley S (1974) *Values of non-atomic games*. Princeton University Press, Princeton
- Azrieli Y, Lehrer E (2007) Market games in large economies with a finite number of types. *Econ Theor* 31:327–342
- Bennett E, Wooders M (1979) Income distribution and firm formation. *J Comp Econ* 3:304–317. <http://www.mymawooders.com/>
- Bergstrom T, Varian HR (1985) When do market games have transferable utility? *J Econ Theor* 35(2):222–233
- Billera LJ (1974) On games without side payments arising from a general class of markets. *J Math Econ* 1(2):129–139
- Billera LJ, Bixby RE (1974) Market representations of n -person games. *Bull Am Math Soc* 80(3):522–526
- Böhm V (1974) The core of an economy with production. *Rev Econ Stud* 41:429–436
- Bondareva O (1963) Some applications of linear programming to the theory of cooperative games. *Problemy kibernetiki* 10 (in Russian, see English translation in *Selected Russian papers in game theory 1959–1965*. Princeton University Press, Princeton)
- Buchanan J (1965) An economic theory of clubs. *Economica* 33:1–14
- Cartwright E, Conley J, Wooders M (2006) The law of demand in Tiebout economies. In: Fischel WA (ed) *The Tiebout model at 50: essays in public economics in honor of Wallace Oates*. Lincoln Institute of Land Policy, Cambridge
- Casella A, Feinstein JS (2002) Public goods in trade on the formation of markets and jurisdictions. *Intern Econ Rev* 43:437–462
- Champsaur P (1975) Competition vs. cooperation. *J Econ Theory* 11:394–417
- Cheng HC (1981) On dual regularity and value convergence theorems. *J Math Econ* 8:37–57
- Conley J, Smith S (2005) Coalitions and clubs; Tiebout equilibrium in large economies. In: Demange G, Wooders M (eds) *Group formation in economies; networks, clubs and coalitions*. Cambridge University Press, Cambridge
- Conley JP, Wooders M (1995) Hedonic independence and taste-homogeneity of optimal jurisdictions in a Tiebout economy with crowding types. *Ann D'Econ Stat* 75(76):198–219
- Crawford VP, Kelso AS (1982) Job matching, coalition formation, and gross substitutes. *Econometrica* 50:1483–1504
- Debreu G, Scarf H (1963) A limit theorem on the core of an economy. *Int Econ Rev* 4:235–246
- Demange G (1994) Intermediate preferences and stable coalition structures. *J Math Econ* 1994:45–48
- Ellickson B, Grodal B, Scotchmer S, Zame W (1999) Clubs and the market. *Econometrica* 67:1185–1218
- Gale D, Shapley LS (1962) College admissions and the stability of marriage. *Am Math Mon* 69:9–15
- Garraff R, Qin C-Z (1997) On a market for coalitions with indivisible agents and lotteries. *J Econ Theor* 77(1):81–101
- Gillies DB (1953) *Some theorems on n -person games*. PhD Dissertation, Department of Mathematics, Princeton University, Princeton

- Haimanko O, Le Breton M, Weber S (2004) Voluntary formation of communities for the provision of public projects. *J Econ Theor* 115:1–34
- Hildenbrand W (1974) Core and equilibria of a large economy. Princeton University Press, Princeton
- Hurwicz L, Uzawa H (1977) Convexity of asymptotic average production possibility sets. In: Arrow KJ, Hurwicz L (eds) *Studies in resource allocation processes*. Cambridge University Press, Cambridge
- Kalai E, Zemel E (1982a) Totally balanced games and games of flow. *Math Oper Res* 7:476–478
- Kalai E, Zemel E (1982b) Generalized network problems yielding totally balanced games. *Oper Res* 30:998–1008
- Kaneko M, Wooders M (1982) Cores of partitioning games. *Math Soc Sci* 3:313–327
- Kaneko M, Wooders M (1986) The core of a game with a continuum of players and finite coalitions; the model and some results. *Math Soc Sci* 12:105–137. <http://www.mymawooders.com/>
- Kaneko M, Wooders M (2004) Utility theories in cooperative games, Chapter 19. In: *Handbook of utility theory*, vol 2. Kluwer, Dordrecht, pp 1065–1098
- Kannai Y (1972) Continuity properties of the core of a market. *Econometrica* 38:791–815
- Konishi H, Le Breton M, Weber S (1998) Equilibrium in a finite local public goods economy. *J Econ Theory*:79224–79244
- Kovalenkov A, Wooders M (2001) Epsilon cores of games with limited side payments: nonemptiness and equal treatment. *Games Econ Behav* 36(2):193–218
- Kovalenkov A, Wooders M (2003) Approximate cores of games and economies with clubs. *J Econ Theory* 110:87–120
- Kovalenkov A, Wooders M (2005) A law of scarcity for games. *Econ Theor* 26:383–396
- Kovalenkov A, Wooders M (2006) Comparative statics and laws of scarcity for games. In: Aliprantis CD, Matzkin RL, McFadden DL, Moore JC, Yannelis NC (eds) *Rationality and equilibrium: a symposium in honour of Marcel K. Richter*, *Studies in economic theory series*, vol 26. Springer, Berlin, pp 141–169
- Mas-Colell A (1975) A further result on the representation of games by markets. *J Econ Theor* 10(1):117–122
- Mas-Colell A (1977) Indivisible commodities and general equilibrium theory. *J Econ Theory* 16(2):443–456
- Mas-Colell A (1979) Competitive and value allocations of large exchange economies. *J Econ Theor* 14:307–310
- Mas-Colell A (1980) Efficiency and decentralization in the pure theory of public goods. *Q J Econ* 94:625–641
- Mas-Colell A (1985) *The theory of general economic equilibrium*. Economic Society Publication No. 9. Cambridge University Press, Cambridge
- Moulin M (1988) Axioms of cooperative decision making. *Econometric Society Monograph No. 15*. Cambridge Press, Cambridge
- Moulin H (1992) Axiomatic cost and surplus sharing, Chapter 6. In: Arrow K, Sen AK, Suzumura K (eds) *Handbook of social choice and welfare*, 1st edn, vol 1. Elsevier, Amsterdam, pp 289–357
- von Neumann J, Morgenstern O (1953) *Theory of games and economic behavior*. Princeton University Press, Princeton
- Owen G (1975) On the core of linear production games. *Math Program* 9:358–370
- Pauly M (1970) Cores and clubs. *Public Choice* 9:53–65
- Qin C-Z, Shapley LS, Shimomura K-I (2006) The Walras core of an economy and its limit theorem. *J Math Econ* 42(2):180–197
- Roth A, Sotomayer M (1990) *Two-sided matching: a study in game-theoretic modeling and analysis*. Cambridge University Press, Cambridge
- Scotchmer S, Wooders M (1988) Monotonicity in games that exhaust gains to scale. IMSSS Technical Report No. 525, Stanford University
- Shapley LS (1952) Notes on the N-Person game III: some variants of the von-Neumann-Morgenstern definition of solution Rand Corporation research memorandum RM-817:1952
- Shapley LS (1964) Values of large games-VII: a general exchange economy with money. Rand Memorandum RM-4248-PR
- Shapley LS (1967) On balanced sets and cores. *Nav Res Logist Q* 9:45–48
- Shapley LS, Shubik M (1960) On the core of an economic system with externalities. *Am Econ Rev* 59:678–684
- Shapley LS, Shubik M (1966) Quasi-cores in a monetary economy with nonconvex preferences. *Econometrica* 34:805–827
- Shapley LS, Shubik M (1969) On market games. *J Econ Theor* 1:9–25
- Shapley LS, Shubik M (1972) The assignment game I; The core. *Int J Game Theor* 1:11–30
- Shapley LS, Shubik M (1975) Competitive outcomes in the cores of market games. *Int J Game Theor* 4:229–237
- Shapley LS, Shubik M (1977) Trade using one commodity as a means of payment. *J Political Econ* 85:937–968
- Shubik M (1959a) Edgeworth market games. In: Luce FR, Tucker AW (eds) *Contributions to the theory of games IV. Annals of mathematical studies* 40. Princeton University Press, Princeton, pp 267–278
- Shubik M (1959b) Edgeworth market games. In: Luce FR, Tucker AW (eds) *Contributions to the theory of games IV. Annals of mathematical studies* 40. Princeton University Press, Princeton, pp 267–278
- Shubik M, Wooders M (1982a) Clubs, markets, and near-market games. In: Wooders M (ed) *Topics in game theory and mathematical economics: essays in honor of Robert J Aumann*. Field Institute Communication Volume, American Mathematical Society, originally Near Markets and Market Games, Cowles Foundation, Discussion Paper No. 657
- Shubik M, Wooders M (1982b) Near markets and market games. Cowles Foundation Discussion Paper No. 657. <http://www.mymawooders.com/>
- Shubik M, Wooders M (1982c) Clubs, markets, and near-market games. In: Wooders M (ed) *Topics in game theory and mathematical economics: essays in honor of Robert J Aumann*. Field Institute Communication

- Volume, American Mathematical Society, originally Near Markets and Market Games, Cowles Foundation, Discussion Paper No. 657
- Shubik M, Wooders M (1983a) Approximate cores of replica games and economies: part II set-up costs and firm formation in coalition production economies. *Math Soc Sci* 6:285–306
- Shubik M, Wooders M (1983b) Approximate cores of replica games and economies: part I replica games, externalities, and approximate cores. *Math Soc Sci* 6:27–48
- Shubik M, Wooders M (1983c) Approximate cores of replica games and economies: part II set-up costs and firm formation in coalition production economies. *Math Soc Sci* 6:285–306
- Shubik M, Wooders M (1986) Near-markets and market-games. *Econ Stud Q* 37:289–299
- Sondermann D (1974) Economics of scale and equilibria in coalition production economies. *J Econ Theor* 8:259–291
- Sun N, Trockel W, Yang Z (2008) Competitive outcomes and endogenous coalition formation in an n-person game. *J Math Econ* 44:853–860
- Tauman Y (1987) The Aumann-Shapley prices: a survey. In: Roth A (ed) *The Shapley value: essays in honor of Lloyd S Shapley*. Cambridge University, Cambridge
- Tauman Y, Urbano A, Watanabe J (1997) A model of multiproduct price competition. *J Econ Theor* 77:377–401
- Tiebout C (1956) A pure theory of local expenditures. *J Political Econ* 64:416–424
- Weber S (1979) On ε -cores of balanced games. *Int J Game Theor* 8:241–250
- Weber S (1981) Some results on the weak core of a non-sidepayment game with infinitely many players. *J Math Econ* 8:101–111
- Winter E, Wooders M (1990) On large games with bounded essential coalition sizes. University of Bonn Sonderforschungsbereich 303 Discussion Paper B-149. <http://www.mynawooders.com/>. *Intern J Econ Theor* (2008) 4:191–206
- Wooders M (1977) Properties of quasi-cores and quasi-equilibria in coalition economies. SUNY-Stony Brook Department of Economics Working Paper No. 184, revised (1979) as A characterization of approximate equilibria and cores in a class of coalition economies. State University of New York Stony Brook Economics Department. <http://www.mynawooders.com/>
- Wooders M (1978) Equilibria, the core, and jurisdiction structures in economies with a local public good. *J Econ Theor* 18:328–348
- Wooders M (1983) The epsilon core of a large replica game. *J Math Econ* 11:277–300. <http://www.mynawooders.com/>
- Wooders M (1988) Large games are market games I. Large finite games. C.O.R.E. Discussion Paper No. 8842. <http://www.mynawooders.com/>
- Wooders M (1989) A Tiebout theorem. *Math Soc Sci* 18:33–55
- Wooders M (1991a) On large games and competitive markets I: theory. University of Bonn Sonderforschungsbereich 303 Discussion Paper No. (B-195, Revised August 1992). <http://www.mynawooders.com/>
- Wooders M (1991b) The efficaciousness of small groups and the approximate core property in games without side payments. University of Bonn Sonderforschungsbereich 303 Discussion Paper No. B-179. <http://www.mynawooders.com/>
- Wooders M (1992a) Inessentiality of large groups and the approximate core property; An equivalence theorem. *Econ Theor* 2:129–147
- Wooders M (1992b) Large games and economies with effective small groups. University of Bonn Sonderforschungsbereich 303 Discussion Paper No. B-215. (Revised: In: Mertens J-F, Sorin S (eds) *Game-theoretic methods in general equilibrium analysis*. Kluwer, Dordrecht). <http://www.mynawooders.com/>
- Wooders M (1993) The attribute core, core convergence, and small group effectiveness; The effects of property rights assignments on the attribute core. University of Toronto Working Paper No. 9304
- Wooders M (1994) Equivalence of games and markets. *Econometrica* 62:1141–1160. <http://www.mynawooders.com/>
- Wooders M (1997) Equivalence of Lindahl equilibria with participation prices and the core. *Econ Theor* 9:113–127
- Wooders M (2007) Core convergence in market games and club economics. *Rev Econ Design* (to appear)
- Wooders M (2008a) Small group effectiveness, per capita boundedness and nonemptiness of approximate cores. *J Math Econ* 44:888–906
- Wooders M (2008b) Games with many players and abstract economies permitting differentiated commodities, clubs, and public goods (submitted)
- Wooders M, Zame WR (1987) Large games; Fair and stable outcomes. *J Econ Theor* 42:59–93
- Zajac E (1972) Some preliminary thoughts on subsidization. Presented at the Conference on Telecommunications Research, Washington

Learning in Games

John Nachbar
Department of Economics, Washington
University, St. Louis, MO, USA

Article Outline

Glossary
Definition of the Subject and Its Importance
Introduction
Deterministic Learning
Stochastic Learning
Future Directions
Bibliography

Glossary

Bayesian learning In repeated games, a model in which each player best responds to her prior, which is a probability distribution over her opponent's behavior strategies.

Behavior strategy In a repeated game, a behavior strategy for player i gives, for each possible date and each possible history of play in the game up to that date, a probability distribution over i 's actions next period. This includes the possibility that player i may play some action for certain.

Belief learning In repeated games, a model in which players best respond to prediction rules.

Prediction rule In a two-player repeated game, a *deterministic prediction rule* gives a probability distribution over the opponent's actions next period as a deterministic function of the history of the game. In a *stochastic* prediction rule, the distribution over the opponent's actions can depend on history probabilistically. A deterministic prediction rule that gives player 2's forecast about player 1's actions is formally equivalent to a behavior strategy for player 1.

Repeated games A repeated game is an extensive form of game representing a repeated strategic interaction. The interaction being repeated is called the *stage game*. In a *discounted* repeated game, payoffs in the repeated game are a geometrically weighted sum of the payoffs each period from the stage game. The weight on period t payoffs is δ^t , where $\delta \in (0, 1)$ is the *discount factor*. A player who is patient has a discount factor close to 1.

Definition of the Subject and Its Importance

In the context of this entry, *learning* refers to a particular class of dynamic game theoretic models. In models in this class, players are rational in the sense that they forecast the future behavior of their opponents and optimize, or ϵ optimize, with respect to their forecasts. But players are not necessarily in equilibrium; in particular, their forecasts are not necessarily accurate. Two objectives are to model out-of-equilibrium behavior by sophisticated players and to understand when, or whether, play might converge to equilibrium.

Learning models are a branch of a larger literature on out-of-equilibrium behavior in dynamic games. In other branches of the literature, players are modeled as "adaptive." Players do not forecast and they do not optimize. Rather, they follow some other form of behavioral rule, such as imitation, regret minimization, or reinforcement. Learning models, especially those that attempt to capture sophisticated behavior, are most appropriate in settings where players have a good understanding of their strategic environment and where the stakes are high enough to make deliberation worthwhile. For surveys of evolutionary/adaptive models, see Sandholm (2007b), Young (2008a, b), and Camerer (2008).

Introduction

This entry surveys research on learning in games. The subject is as old as game theory itself, dating to the work of Cournot in the early nineteenth century. But, with some important exceptions, research on learning in games was largely dormant until the late 1980s.

I have divided the survey into two parts. The first part covers *deterministic* learning, in which each player optimizes with respect to a *deterministic prediction rule*, which gives a probability distribution over the opponent's next period actions as a deterministic function of the history of the game to date. Cournot's original dynamic model falls into this category. The second part briefly surveys some of the literature on *stochastic* learning, in which each player optimizes with respect to a prediction rule that is a probabilistic function of history. This survey expands and updates my earlier survey, Nachbar (2008). For a complementary survey, with a different emphasis and more detail on some topics, see Fudenberg and Levine (2009).

One can formulate learning within the context of very general dynamic games. For this survey, however, I focus on the most studied case: discounted infinitely repeated games with perfect monitoring (at each stage, players know the sequence of action profiles chosen at previous stages). In section “[Convergence to What Sort of Equilibrium?](#),” I discuss an important variant: repeated extensive form games in which the perfect monitoring assumption is violated. Also for simplicity, I focus mostly on two-player games.

Deterministic Learning

Classical Learning Models

I begin by introducing two classical learning models, Cournot best response dynamics and fictitious play. In both models, players are unsophisticated: they cannot learn to forecast even simple patterns in opponent behavior, such as deterministic alternation between two stage-game actions. But Cournot best response and fictitious play are historically important, easy to describe, relatively tractable, heavily studied, and often used as building blocks in more elaborate dynamic models

(e.g., Sandholm 2007a) and can be used to illustrate some of the major themes in the literature.

Cournot Best Response Dynamics

Following Cournot (1838), consider two mineral spring owners, 1 and 2, who are each trying to make money by selling water. Suppose that the profit to firm 1 is $(13 - Q^1 - Q^2)Q^1 - Q^1$. Here, Q^i is the output in one period by firm i , $13 - Q^1 - Q^2$ gives the market price, and output costs 1 per unit; hence total cost to firm 1 is Q^1 . The quantity Q^1 that maximizes firm 1's profit depends on Q^2 , and so in order to choose Q^1 , firm 1 must forecast Q^2 . In a model of sophisticated behavior, firm 1 would want to forecast not only Q^2 next period but also how firm 2 might respond in subsequent periods. But, for the moment, assume that firm 1 is myopic, meaning that firm 1 chooses Q^1 to maximize next period's expected profit. Then firm 1 should set,

$$Q^1 = \frac{12 - \mathbb{E}[Q^2]}{2},$$

where $\mathbb{E}[Q^2]$ is firm 1's expectation of Q^2 .

What should $\mathbb{E}[Q^2]$ be? In Cournot's story, firm 1 puts probability one on firm 2 continuing to produce next period whatever firm 2 produced this period, hence $\mathbb{E}[Q_{t+1}^2] = Q_t^2$. Cournot argued that if both firms adjust in this way, then eventually Q^1 and Q^2 converge to the same number, Q^* , which must satisfy,

$$Q^* = \frac{12 - Q^*}{2},$$

hence

$$Q^* = 4.$$

In modern terminology, play in this model of out-of-equilibrium learning eventually converges to the symmetric Nash equilibrium of the single-period version of this game, namely, (4, 4). Cournot thus not only identified what is now called the Nash equilibrium of this game, but he also provided a dynamic rationale for why the Nash equilibrium might arise.

Fictitious Play

Fictitious play and its variants, first introduced by Brown (1951), is the single most heavily studied learning model. The term “fictitious play” refers to the idea that players would run through the fictitious play dynamic in their minds, prior to the start of actual play, as a way of forecasting their opponent’s behavior once play begins. In recent years, however, fictitious play has usually been interpreted as a model of actual behavior in real time.

In its modern form, fictitious play dictates that each player predicts that the probability that her opponent will play an action, say L , next period is a weighted sum of an initial probability on L and the frequency with which L has been chosen to date. The weight on the frequency is $t/(t + k)$, where t is the number of periods thus far and $k > 0$ is a fixed parameter. The larger is k , the longer the initial probability on L dominates the player’s forecast about L .

In fictitious play, players are *as if* myopic: for any k and any discount factor, each player effectively believes he cannot influence his opponent’s future behavior (I discuss this further in section “Belief Learning and Bayesian Learning”), and so it is optimal to maximize current period payoff, ignoring the future. As this observation suggests, behavior under fictitious play is naive. Fictitious play is, in fact, mathematically very close to regret minimization, an adaptive model; see Hart and Mas-Colell (2000, 2001). To justify fictitious play for sophisticated players, Fudenberg and Kreps (1993) suggest thinking of environments in which there is a large population of players randomly matched (an environment normally associated with evolutionary models), rather than a standard repeated game; see also Fudenberg and Takahashi (2011).

Belief Learning and Bayesian Learning

In general, *belief learning* refers to any learning model in which players best respond (or ε -best respond) to a deterministic prediction rule. Both the Cournot dynamic and fictitious play are examples of belief learning. An alternative approach to modeling learning is *Bayesian*. Recall that in a repeated game, a *behavior strategy* gives, for every history, a probability over the player’s

stage-game actions next period. In a Bayesian learning model, each player chooses a behavior strategy that best responds (or ε -best responds) to a prior, a probability distribution over the opponent’s behavior strategies. A basic observation in the theory of deterministic learning is that belief learning and Bayesian learning are mathematically equivalent.

Explicitly, player 1’s prediction rule about player 2 is mathematically identical to a behavior strategy for player 2. Thus, any belief learning model is equivalent to a Bayesian model in which each player optimizes with respect to a prior that places probability one on her prediction rule, now reinterpreted as the opponent’s behavior strategy.

Conversely, any Bayesian model is equivalent to a belief learning model. Explicitly, for any prior over player 2’s behavior strategies, there is a degenerate prior, assigning probability one to a particular behavior strategy, that is, equivalent in the sense that both priors induce the same distributions over play in the game, no matter what behavior strategy player 1 herself adopts. This is a form of Kuhn’s Theorem, Kuhn (1964); for a version applicable to repeated games, see Aumann (1964). I refer to the behavior strategy used in the degenerate prior as a *reduced form* of the original prior. Thus, any Bayesian model is equivalent to a Bayesian model in which each player’s prior places probability one on the reduced form, and any such Bayesian model is equivalent to a belief learning model.

To illustrate, consider fictitious play. For simplicity, consider a stage game with just two actions, L and R . By an i.i.d. strategy for player 2, I mean a behavior strategy in which player 2 plays L with probability q , independent of history. Thus, if $q = 1/2$, then player 2 always randomizes 50:50 between L and R . Fictitious play is equivalent to a degenerate Bayesian model in which each player places probability one on the fictitious play prediction rule, and one can show that this is equivalent in turn to a nondegenerate Bayesian model in which the belief is represented as a beta distribution over q ; see, for example, Fudenberg and Levine (1998). For example, the uniform distribution over q corresponds to taking

the initial probability of L to be $1/2$ and the parameter k to be 2. Thus, fictitious play can be thought of as a Bayesian model in which each player is certain that the other is playing an i.i.d. strategy, but isn't sure which i.i.d. strategy. Note that since each player believes the other is i.i.d., each believes she has no influence on the other's future behavior. Hence, as noted earlier, myopic optimization is optimal for any discount factor.

Likewise, the Cournot prediction rule is Bayesian. In particular, one can form the degenerate prior from the Cournot prediction rule. To my knowledge, the Cournot rule has no compelling nondegenerate Bayesian interpretation.

The fact that both the Cournot dynamic and fictitious play can be interpreted as Bayesian underscores the fact that there is no presumption that Bayesian players are sophisticated. I take up the issue of sophisticated learning below, especially in section "[Sophisticated Learning](#)."

What Should "Convergence" Mean?

One of the characteristics of the literature on learning in games is a running dialog about what "converges to equilibrium" ought to mean.

Convergence in What Game?

For concreteness, suppose that we are looking for convergence to Nash equilibrium; similar comments apply to convergence to other forms of equilibrium, such as correlated equilibrium.

In the models that are the focus of this entry, learning takes place within the context of a repeated game. Therefore, it is natural to look for convergence to Nash equilibria of the repeated game. Repeated play of a stage-game Nash equilibrium always constitutes a Nash equilibrium of the repeated game, but typically there are also other types of repeated game equilibria.

First, even if players do play stage-game Nash equilibria, they need not play the *same* stage-game Nash equilibrium in every period. As an example, consider the repeated battle of the sexes, a stage game of which is given in Fig. 1. Alternation between (a, a) in odd periods and (b, b) in even periods, both of which are stage-game Nash equilibria, is a natural repeated game Nash equilibrium.

Second, if the discount factor is close enough to 1, the folk theorem implies that there will typically be many repeated game Nash equilibrium and some of these may involve play of actions that are not part of any stage-game equilibrium (for more on the folk theorem, see Fudenberg and Tirole (1991)). The standard example is the repeated prisoner's dilemma, in which the repeated game can have Nash equilibria that sustain cooperation, even though cooperation is strictly dominated in the stage game.

More generally, learning stories take place in the context of a larger dynamic game. Learning may lead to Nash equilibrium play in the continuation of this larger game, but Nash equilibrium in the continuation of the larger game need not correspond to Nash equilibrium play in a component (such as a stage game) of that larger game.

Convergence in What Sense?

Until the 1990s, most of the work on learning focused on whether the empirical marginal frequencies of realized play converged to a Nash equilibrium of the stage game.

As an illustration, consider repeated matching pennies, the stage game for which is given in Fig. 2. The Nash equilibrium of repeated matching pennies, for any discount factor, calls for players to randomize 50:50 in every period, regardless of history. Under convergence of the empirical marginal frequencies, play is said to converge to Nash equilibrium if each player plays a half of the time. This criterion is satisfied if play alternates deterministically between (a, a) and (b, b) . In contrast, a Nash equilibrium play path looks random; in particular, with high probability, over any finite set of dates, the four pure

Learning in Games,

Fig. 1 Battle of the sexes

	a	b
a	8, 10	0, 0
b	0, 0	10, 8

Learning in Games,

Fig. 2 Matching pennies

	a	b
a	1, -1	-1, 1
b	-1, 1	1, -1

outcomes (a, a) , (a, b) , (b, a) , and (b, b) occur about equally often. Convergence of the empirical marginal frequencies is thus a very weak convergence criterion.

A somewhat tougher criterion is convergence of the empirical *joint* frequencies. In the matching pennies example, this would require that each of the four pure action profiles gets played 1/4 of the time. This convergence criterion eliminates the previous example but it is still extremely weak. It would, for example, allow convergence to the deterministic sequence $(a, a), (a, b), (b, a), (b, b), (a, a), \dots$ to count as convergence to a Nash equilibrium. Note that along this sequence, (a, a) gets played for certain at dates $t = 1, 5, 9, \dots$. In contrast, in the Nash equilibrium, (a, a) gets played only 1/4 of the time at those dates. These difficulties with defining convergence in terms of empirical frequencies were first emphasized in Fudenberg and Kreps (1993); see also the related Jordan (1993).

Choice of convergence standard is thus a matter of balance. Choose too weak a standard, as in the examples above, and convergence is arguably not meaningful. Choose too strong a standard and one gets impossibility results, even though positive results are still available for other, weaker but arguably still satisfactory, forms of convergence. For a stark example, consider the battle of the sexes game of Fig. 1. If players start in the repeated game Nash equilibrium in which they play (a, a) in odd periods and (b, b) in even periods then play converges (trivially) to Nash equilibrium play in any continuation game. But note that, strictly speaking, we get a *different* Nash equilibrium depending on whether the starting date of the continuation game is odd or even. So a convergence standard that requires that play be close to the *same* repeated game Nash equilibrium in every continuation game yields non-convergence, even in this trivial example. For a more subtle example illustrating the trade-offs in choosing the right convergence standard, see section “[Payoff Uncertainty](#).”

Loosely, the strongest form of convergence that one can generally hope for is that, to an outside observer, play over finite continuation histories looks asymptotically like play of some Nash equilibrium or ε -Nash equilibrium of the repeated game.

Convergence to What Sort of Equilibrium?

Although my focus here is primarily on Nash equilibrium, other solution concepts are often of interest. Obvious alternatives are rationalizable strategy profiles and correlated equilibria, and both have received attention (e.g., Bernheim 1984; Foster and Vohra 1997; Nyarko 1994).

Two other variants of Nash equilibrium have also proved important in the literature. First, it is often natural to assume that players ε optimize rather than exactly optimize. If players only ε optimize then convergence will be to an ε -Nash equilibrium rather than to an exact Nash equilibrium. It is common in the literature to model myopic players as using a logit selection from the stage game’s ε -best response correspondence. In this case, ε -Nash equilibrium in the stage game takes the form of a *quantal response equilibrium* (QRE), McKelvey and Palfrey (1995), a solution concept that has become important in the experimental game theory literature (see, for example, Goeree and Holt 2001). While any Nash equilibrium is an ε -Nash equilibrium, ε -Nash equilibria are typically not Nash equilibria, although the difference is small if ε is small.

A second important variant arises in repeated games in which the stage game has a nontrivial extensive form. In such settings, the perfect monitoring assumption may be untenable. It may make sense to assume instead that while players observe the outcome of the stage game, they do not observe the full strategy profile for that stage game.

To take a concrete example, consider the repeated ultimatum bargaining game. In the stage game, player 1 makes an offer in the form of an integer $x \in [0, 100]$ and player 2 either rejects, yielding a payoff profile of $(0, 0)$, or accepts, yielding a payoff profile of $(x, 100 - x)$. It may make sense to assume that while player 1 can observe player 2’s action (accept or reject) in response to the actual offer, she cannot observe player 2’s entire stage-game strategy, since this would mean observing how player 2 would have responded to every other possible offer.

For this sort of setting, Fudenberg and Levine (1993a) propose *self-confirming equilibrium* (SCE) as an alternative to Nash equilibrium.

There are actually several forms of SCE, depending on the setting and the information of the players. The basic idea is that in a SCE of the stage game, players are optimizing with respect to predictions that are correct along the play path within the stage game. Predictions about how the opponent would behave off the play path, however, may be wrong. SCE bears a family resemblance to conjectural equilibrium (Hahn 1977) and subjective equilibrium (Kalai and Lehrer 1993b).

For players to learn to make correct predictions about behavior off the play path within the stage game, they must experiment with alternative play. Fudenberg and Levine (1993b) show, in the context of a large population steady-state learning model, that if players are sufficiently patient, then they have enough incentive to experiment, learn to make correct predictions, and learn to play a Nash equilibrium of the stage game. For a provocative application of this learning model, see Fudenberg and Levine (2006).

Finally, some other solution concepts have received attention, notably persistent retracts, Kalai and Samet (1984), CURB (closed under rational behavior) sets, Basu and Weibull (1991), and minimal prep sets, Voorneveld (2004, 2005). These solution concepts tend to arise most readily in stochastic learning models, which I discuss in section “Stochastic Learning.”

Convergence in Classical Learning Models

Neither Cournot best response dynamics nor fictitious play exhibits universal convergence, by which I mean convergence for all stage games. It is easy to construct non-convergence examples for Cournot best response, even for games with pure-strategy equilibria. For fictitious play, there are classes of games that cause convergence problems even when one considers only weak forms of convergence, such as convergence of the empirical marginal distributions. The classic example of a problem game, from Shapley (1962), is given in Fig. 3. In this game, if the prediction rules concentrate initial probability on, say, (b, a) , then play cycles, starting at (a, a) then moving to (a, c) , then to (c, c) , and so on. Moreover the cycle becomes slower and slower, so that not even the empirical

Learning in Games,

Fig. 3 The Shapley game

	<i>a</i>	<i>b</i>	<i>c</i>
<i>a</i>	1, 0	0, 0	0, 1
<i>b</i>	0, 1	1, 0	0, 0
<i>c</i>	0, 0	0, 1	1, 0

marginal frequencies converge to the unique Nash equilibrium of the stage game, which has each player randomize equally over a , b , and c . See also Jordan (1993). Hart and Mas-Colell (2003) prove a kind of generalization of the Shapley example: for a large class of dynamics that include a continuous time version of fictitious play as a special case, one can always find a stage game and initial conditions for which convergence to stage-game equilibrium play fails.

On a more positive note, Cournot best response dynamics do converge to pure-strategy stage-game equilibrium if the stage game is solvable by the iterated deletion of strictly dominated strategies. See Bernheim (1984) and Moulin (1984) and, for a more general class of models, Milgrom and Roberts (1991).

As for fictitious play, there are ε -optimizing variants of fictitious play that yield convergence to approximate stage-game equilibrium play (possibly mixed) for all zero-sum games, all games with an interior ESS (evolutionarily stable strategy), and all common interest games, in addition to all games that are strict dominance solvable, with the approximation closer the smaller is ε . Somewhat weaker convergence results are available for supermodular games. These claims follow from results in Hofbauer and Sandholm (2002), which builds on Fudenberg and Kreps (1993) and Benaïm and Hirsch (1999) and which also provides additional references to the large literature on fictitious play.

The use of ε optimization, or something analogous (such as small, privately observed, payoff shocks in each period), is necessary. In a standard repeated game setting, for almost every stage game, standard fictitious play, with exact optimization, yields pure continuation play paths (see Fudenberg and Kreps 1993). Thus, in games like repeated matching pennies, standard fictitious play cannot give convergence to an equilibrium play path. This is an instance of a recurring

modeling theme in the learning literature: it is often easier to get positive convergence results for ε optimization than for exact optimization.

Kalai-Lehrer Learning

Kalai and Lehrer (1993a) (hereafter KL) take a Bayesian perspective and ask what conditions on priors are sufficient to give convergence to equilibrium, or approximate equilibrium, play. I find it helpful to characterize KL, and related papers, in the following way.

A player *learns to predict the play path* if her prediction of next period's play is asymptotically as good as if she knew her opponent's behavior strategy. If the behavior strategies call for randomization then players accurately predict the distribution over next period's play rather than the realization of next period's play. For example, consider a 2×2 game in which player 1 has stage-game actions T and B and player 2 has stage-game actions L and R . If player 2 is randomizing 50:50 every period and player 1 learns to predict the play path, then for every $\gamma > 0$, there is a time, which depends on the realization of player 2's strategy, after which player 1's next period forecast puts the probability of L within γ of $1/2$. (This statement applies to a set of play paths that arises with probability one with respect to the underlying probability model; I gloss over this sort of complication both here and below.) Finally, say that a prior profile (giving a prior for each player) has the *learnable best response property* (LBR) if there is a profile of best response (or ε -best response) strategies (LBR strategies) such that, if the LBR strategies are played, then each player learns to predict the play path.

If LBR holds, and players are using their LBR strategies, then, asymptotically, the continuation play path is an approximate equilibrium play path of the continuation repeated game. The exact sense in which play converges to equilibrium play depends on the strength of learning and of optimization. See KL and also Sandroni (1998) (both of which focus on exact optimization) and Noguchi (2015b) (which considers ε optimization).

KL, building on work in the probability literature on the merging of measures, notably Blackwell and Dubins (1962), show that a strong form

of LBR holds if beliefs satisfy an absolute continuity condition: each player assigns positive probability to any (measurable) set of play paths that has positive probability given the players' actual strategies. A strong sufficient condition for this is that each player assigns positive, even if extremely low, prior probability to her opponent's actual strategy, a condition that KL call *grain of truth*.

Universal Convergence

The work described thus far leaves open whether one can write down *any* deterministic learning model, even an implausible one, that exhibits universal convergence, that is, a learning model that, for a given stage-game form (giving stage-game actions but not payoff functions), gives convergence to Nash equilibrium play in the repeated game for all (or at least for generic) discount factors and stage-game payoff functions.

Implicit in most work on this question is an assumption that the learning model is "uncoupled" in the sense of Hart and Mas-Colell (2003): player i 's prior over her opponent's repeated game strategies does not depend on the specification of her opponent's stage-game payoffs. Cournot best response and fictitious play, as typically employed, are both examples of uncoupled learning models. One can get convergence in a "coupled" model simply by having the players play a Nash equilibrium of the repeated game. So, some degree of uncoupling is needed to avoid triviality. On the other hand, full uncoupling is a strong assumption from the perspective of sophisticated learning, since it effectively rules out introspective reasoning about one's opponent. And full uncoupling presumably makes convergence more difficult to achieve. Indeed, the main result of Hart and Mas-Colell (2003) is an impossibility theorem on convergence for certain classes of adaptive learning models.

For belief learning models, there are two competing intuitions on (uncoupled) universal convergence. One intuition is that universal convergence, even to ε -equilibrium, is impossible because universal learnability is impossible: for any prior, there are opposing strategies that a player will fail to forecast, even approximately. One can

show this via a diagonalization argument along the lines of Oakes (1985), and it also follows as a corollary of results discussed in section “[Sophisticated Learning](#).” In fact, for any given prior, the set of strategies one can learn to forecast is small, in a sense that can be made precise. A classic reference on the difficulty of prediction, in the context of general stochastic processes, is Freedman (1965). For a useful survey, see the literature review in Al-Najjar (2009).

The competing intuition is that universal learnability is not necessary for convergence. The fact that both players are engaged in a game provides structure that conceivably could force posteriors to be correct, at least along the path of play, even if priors are fundamentally wrong. Something like this happens in fictitious play. In the standard Bayesian interpretation of fictitious play (section “[Belief Learning and Bayesian Learning](#)”), each player is certain that her opponent is playing an i.i.d. strategy, even though neither player is (since, for many games, an i.i.d. strategy is not optimal under the fictitious play prediction rule). Yet, in many games, fictitious play generates behavior that looks asymptotically i.i.d. and play does converge to that of a Nash equilibrium.

Based on the first intuition and also on other negative results like those in Hart and Mas-Colell (2003) and Foster and Young (2003), it had been widely believed that universal convergence was impossible for deterministic learning. More recently, however, Noguchi (2015a) has reported a universal convergence result for ε -equilibrium.

Sophisticated Learning

A number of papers investigate classes of prediction rules that exhibit desirable properties, such as the ability to detect certain kinds of patterns in opponent behavior. Important examples include Aoyagi (1996), Fudenberg and Levine (1995), Fudenberg and Levine (1999), and Sandroni (2000).

In Nachbar (2005), I consider the issue of sophistication from a Bayesian perspective. For simplicity, focus on two-player games. Fix a profile of priors and a subset of behavior strategies for each player, and consider the following criteria for these strategy subsets.

- *Learnability.* For any strategy profile drawn from the strategy subsets, both players learn to predict the play path.
- *Richness.* If a behavior strategy is included in one of the strategy subsets then (informally) certain variations on that strategy must be included in *both* strategy subsets. This condition, called CSP in Nachbar (2005), is satisfied automatically if the strategy subsets consist of all strategies satisfying a standard complexity bound, the same bound for both players. Thus richness/CSP holds if the subsets consist of all strategies with k -period memory, or all strategies that are automaton implementable, or all strategies that are Turing implementable, and so on.
- *Consistency.* Each player’s subset contains a best response to her belief. The motivating idea is that, for priors to be considered sophisticated, a necessary (but not sufficient) condition is that the priors can be represented as probability distributions whose supports (informally speaking) contain sets satisfying these criteria. Note that typically priors have multiple representations. For example, as discussed earlier, the prior for fictitious play could be represented either as a beta distribution over i.i.d. strategies or as a degenerate prior that assigns probability one to the associated reduced form. The question, therefore, is not whether all representations of a given prior satisfy the above criteria but whether any do. Nachbar (2005) studies this question and concludes that in a large number of cases, there is no representation of any prior that satisfies these criteria.

Consider, for example, the Bayesian interpretation of fictitious play in which priors are probability distributions over the i.i.d. strategies. The set of i.i.d. strategies satisfies learnability and richness. But for any stage game in which neither player has a weakly dominant action, the i.i.d. strategies violate consistency: any player who is optimizing will not be playing i.i.d. The main result in Nachbar (2005) implies that any other interpretation of fictitious play will violate at least one of the above criteria.

As shown in Nachbar (2005), this feature of Bayesian fictitious play extends to all Bayesian learning models. For large classes of repeated games, for *any* profile of priors, and for *any* strategy subsets satisfying richness, if learnability holds, then consistency fails.

Let me make a few remarks. First, since the set of all strategies satisfies richness and consistency, it follows that for any profile of priors there is a strategy profile that the players will not learn to predict. This can also be shown directly by a diagonalization argument along the lines of Oakes (1985) and Dawid (1985). The impossibility result of Nachbar (2005) can be viewed as a game theoretic version of Dawid (1985). For a description of what subsets *are* learnable, see Noguchi (2015b).

Second, suppose that the strategy subsets are generated by some standard definition of complexity, the same for both players. Then, as noted above, richness holds. Suppose further that there are priors for which learnability holds. This will be the case, for example, for Turing implementable strategies, since the set of such strategies is countable. Then, for such priors, for large classes of repeated games, consistency fails: best responses must violate the complexity bound.

Third, if one constructs a Bayesian learning model satisfying learnability and consistency, then LBR (see section “[Kalai-Lehrer Learning](#)”) holds, and, if players play their LBR strategies, play converges to equilibrium play. This identifies a potentially attractive class of Bayesian models in which convergence obtains. The impossibility result says, however, that if learnability and consistency hold, then player beliefs must be partially equilibrated in the sense of, in effect, excluding some of the strategies required by richness.

Fourth, the main result in Nachbar (2005) is robust along a number of dimensions. It holds under ε optimization, for ε small. It holds for fairly weak definitions of prediction (e.g., definitions that allow occasional but persistent forecasting errors). And the learnability condition can be relaxed. A weaker learnability condition would require that a player be required to learn to predict only when her own strategy is optimal. A variation of the result states that, for a large class of repeated

games, if, say, player 1’s strategy is (ε) optimal and drawn from her strategy subset, and if those subsets satisfy the richness condition, then there is a strategy in the opponent’s subset for which player 1 cannot learn to predict the play path.

Last, consistency is not *necessary* for convergence. See the discussion of fictitious play in section “[Universal Convergence](#).” The impossibility result is a statement about the ability to construct Bayesian models with certain properties; it is not a statement about convergence to equilibrium *per se*.

Payoff Uncertainty

Suppose that, at the start of the repeated game, each player is privately informed of his or her stage-game payoff function, which remains fixed throughout the course of the repeated game. Refer to player i ’s stage-game payoff function as her *payoff type*. Assume that the joint distribution over payoff functions is independent (to avoid correlation issues that are not central to my discussion) and commonly known.

Each player can condition her behavior strategy in the repeated game on her realized payoff type. A mathematically correct way of representing this conditioning is via distributional strategies; see Milgrom and Weber (1985).

For any prior about player 2, now a probability distribution over player 2’s distributional strategies, and given the probability distribution over player 2’s payoff types, there is a behavior strategy for player 2 in the repeated game that is equivalent in the sense that it generates the same distribution over play paths. Again, this is essentially Kuhn’s theorem.

Say that a player *learns to predict the play path* if her forecast of next period’s play is asymptotically as good as if she knew the reduced form of her opponent’s distributional strategy. This definition specializes to the previous one if the distribution over types is degenerate. If distributional strategies are in Nash equilibrium (also known in this context as a Bayesian Nash equilibrium), then, in effect, each player is optimizing with respect to a degenerate belief that puts probability one on her opponent’s actual distributional

strategy, and in this case players trivially learn to predict the path of play.

One can define LBR (see section “[Kalai-Lehrer Learning](#)”) for distributional strategies and, as in the payoff certainty case, one can show that LBR implies convergence to Nash equilibrium play in the repeated game with payoff types.

More interestingly, Nash equilibrium play in the repeated game with payoff types implies convergence to Nash equilibrium play of the *realized* repeated game – the repeated game determined by the realized type profile. This line of research was initiated by Jordan (1991). Other important papers include Kalai and Lehrer (1993a) (KL), Jordan (1995), Nyarko (1998), and Jackson and Kalai (1999) (the last studies recurring rather than repeated games).

Suppose first that the realized type profile has positive probability. In this case, if a player learns to predict the play path then, as shown by KL, her forecast is asymptotically as good as if she knew both her opponent’s distributional strategy *and* her opponent’s realized type. LBR then implies that actual play, meaning the distribution over play paths generated by the realized behavior strategies, converges to equilibrium play of the realized repeated game. For example, suppose that the type profile for matching pennies gets positive probability. In the unique equilibrium of repeated matching pennies, players randomize 50:50 in every period. Therefore, LBR implies that if the matching pennies type profile is realized then each player’s behavior strategy in the realized repeated game involves 50:50 randomization asymptotically.

If the distribution over types admits a density, so that no type profile receives positive probability, then convergence is more complicated. Suppose that players are myopic and that the realized stage game is like matching pennies, with a unique and fully mixed equilibrium. Given myopia, the unique equilibrium of the realized repeated game calls for repeated play of the stage-game equilibrium. In particular, it calls for players to randomize. It is not hard to show, however, that in a type space game with a density, exact optimization calls for each player to play a pure strategy for almost every realized type (this is a generalization of a point made in the context of fictitious play in section

“[Convergence in Classical Learning Models](#)”). Thus, for almost every realized type profile in a neighborhood of a game like matching pennies, actual play (again meaning the distribution over play paths generated by the realized behavior strategies) cannot converge to Nash equilibrium play of the realized repeated game, *even if the distributional strategies are in Nash equilibrium*. Foster and Young (2001) provide a generalization for non-myopic players.

This impossibility result is not robust to weakening optimization to ε optimization; the positive convergence results for ε -optimizing fictitious play provide one illustration; see section “[Convergence in Classical Learning Models](#).” A more subtle point, however, is that, even for exact optimization, a form of convergence obtains that, while weaker than convergence of actual play, is still very strong.

For simplicity, assume that each player knows the other’s distributional strategy and that these strategies form a (Bayesian) Nash equilibrium. Then to an outsider, for almost any type profile, observed play looks asymptotically like Nash equilibrium play in the realized repeated game (this follows from the main theorem in Nyarko (1998)). In particular, in a neighborhood of a game like matching pennies, for almost any type profile, observed play looks random. Since, by Foster and Young (2001), actual play in this setting cannot converge to equilibrium and in particular cannot be random, the implication is that convergence to equilibrium involves a form of purification in the sense of Harsanyi (1973), a point that has been emphasized by Nyarko (1998) and Jackson and Kalai (1999). To a player in the game, opponent behavior likewise looks random because, even if she knows her opponent’s distributional strategy, she does not know her opponent’s type. As play proceeds, each player in effect learns more about her opponent’s type, but never enough to zero in on her opponent’s realized, pure, behavior strategy.

Finally, the difficulties with characterizing sophistication in Bayesian learning models extend to models with payoff uncertainty, with learnability, richness, and consistency redefined in terms of distributional strategies; see Nachbar (2001).

Stochastic Learning

For a concrete example of a stochastic prediction rule, consider Young (1993), which uses a variant of fictitious play. In standard fictitious play, a player gets to observe the entire history of the game to date. In a bounded memory version of fictitious play, a player gets to observe only the last, say, s periods. In Young's variation on bounded memory fictitious play, a player observes only an $I! < s$ -period sample of the last s periods, with the sample drawn according to a probability distribution that assigns positive probability to every possible sample of length $I!$. If the player's forecast about period $t + 1$ depends on her information in period t , then her prediction rule will be stochastic: her forecast will depend on the realization of the date t sample. More generally, a player with a (nondegenerate) stochastic prediction rule will typically not be able to compute what her own date t forecast will be conditional on the date t history. This implies that a player with a (nondegenerate) stochastic prediction rule is not Bayesian.

A range of stochastic learning models have been studied, nearly all for myopic optimization. In some cases (e.g., Foster and Young 2003), the stochastic prediction rules have a quasi-Bayesian interpretation: most of the time, players optimize with respect to fixed prediction rules, as in a Bayesian model, but occasionally players switch to new prediction rules, implicitly abandoning their priors.

In this last part of the survey, I discuss two of the main themes in the stochastic learning literature. One is familiar: convergence to equilibrium. The other is new: calibration.

Convergence in Stochastic Learning Models

Although details can differ substantially from model to model, convergence results for stochastic learning models often have a similar feel. The stochastic nature of the prediction rule introduces an element of randomness into behavior. This randomness can cause players to stumble upon close-to-equilibrium play in the stage game and players, once near equilibrium, tend to stay near equilibrium, if not forever then at least for long periods of time.

One thread of the literature on convergence in stochastic learning models poses specific models and investigates what sort of solution concepts emerge. In the case of the Young (1993) stochastic fictitious play model, the stochastic dynamics serve as a selection mechanism: learning tends to concentrate on some stage-game Nash equilibrium but not others, even distinguishing between multiple strict Nash equilibrium, a phenomenon that also appears in many stochastic evolutionary models. In other stochastic learning models, convergence is sometimes not to Nash equilibrium but instead to sets containing Nash equilibria, such as CURB sets or minimal prep sets; see Hurkens (1995), Sanchirico (1996), and Kets and Voorneveld (2008).

A second thread of the literature asks whether it is possible to construct *any* stochastic learning model that yields universal convergence to Nash equilibrium. The focus on stochastic prediction rules was motivated in part by the failure to find universal convergence results using deterministic prediction rules; see also section "[Universal Convergence](#)." The seminal paper on global convergence with stochastic prediction rules is Foster and Young (2003). See also Hart and Mas-Colell (2006) and Germano and Lugosi (2007).

Calibration

Foster and Vohra (1998) showed that if one allows stochastic prediction rules then one can get approximate calibration on *every* realized history. That is, no matter what strategy the opponent plays, the prediction rule will approximately satisfy a statistical condition called calibration – and one can make the approximation to exact calibration as close as one likes. Calibration, in turn, implies a weak form of convergence to correlated equilibrium; see Foster and Vohra (1997).

Calibration is a weak condition, so weak that it is possible for a prediction rule to be calibrated without it being accurate in any reasonable sense; see Kalai et al. (1999). Nevertheless, calibration is sufficiently strong that a simple diagonalization argument shows that no deterministic prediction rule can satisfy it even approximately; see Oakes (1985). The discovery that one can get approximate calibration using stochastic prediction was, therefore, startling.

The Foster and Vohra (1998) calibration result has spawned another literature, distinct from its application to learning, on whether it is possible for an observer to distinguish between an expert (someone who knows the true form of some stochastic process) and a charlatan (someone who uses the Foster and Vohra (1998) prediction rule, or one of its more elaborate cousins, to pass the observer's tests). Papers in this area include Lehrer (2001), Sandroni (2003), Dekel and Feinberg (2006), Al-Najjar and Weinstein (2008), and Feinberg and Stewart (2008).

Future Directions

Directions for future work, both for the learning literature and, more broadly, for the general literature on out-of-equilibrium dynamics in games, include the following:

1. *Benchmark learning.* One of the goals of the literature is to establish a benchmark model of sophisticated learning. Nachbar (2005) sheds doubt on whether this is possible, but the desiderata considered in Nachbar (2005) are not based on axiomatic, decision theoretic criteria.
2. *Environments.* This survey has focused on repeated two-player games, but many other learning environments and patterns of interaction are possible. One can, for example, study environments in which players are linked through a network as in Blume (1993), or environments in which players are behaviorally heterogeneous. The possibilities are endless. The task is to enrich the scope of research without simply generating an ever expanding catalog of models.
3. *Empirical testing.* One objective of the learning literature is to understand how real people behave in real games. For a survey of work tying learning models to actual behavior see Camerer (2008). The appropriateness of a model may depend on the environment, and in at least in some cases, behavioral heterogeneity may be relevant.
4. *Application.* An understanding of how players behave outside of equilibrium will enrich the application of game theoretic models. One

example concerns institutional design, an instance of which is the choice of game for selling goods (a particular type of auction, for instance). Different institutions may have similar Nash equilibria but very different dynamic properties, and the dynamic properties could therefore play a role in the choice of institution.

Bibliography

- Al-Najjar N (2009) Decision makers as statisticians. *Econometrica* 77(5):1339–1369
- Al-Najjar N, Weinstein J (2008) Comparative testing of experts. *Econometrica* 76(3):541–559
- Aoyagi M (1996) Evolution of beliefs and the Nash equilibrium of normal form games. *J Econ Theory* 70:444–469
- Aumann R (1964) Mixed and behaviour strategies in infinite extensive games. In: Dresher M, Shapley LS, Tucker AW (eds) *Advances in game theory*, *Annals of mathematics studies*, vol 52. Princeton University Press, Princeton, pp 627–650
- Basu K, Weibull AW (1991) Strategy subsets closed under rational behavior. *Econ Lett* 36:141–146
- Benaïm M, Hirsch M (1999) Mixed equilibria arising from fictitious play in perturbed games. *Games Econ Behav* 29:36–72
- Bernheim BD (1984) Rationalizable strategic behavior. *Econometrica* 52(4):1007–1028
- Blackwell D, Dubins L (1962) Merging of opinions with increasing information. *Ann Math Stat* 38:882–886
- Blume LE (1993) The statistical mechanics of strategic interaction. *Games Econ Behav* 5(3):387–424
- Brown GW (1951) Iterative solutions of games by fictitious play. In: Koopmans TJ (ed) *Activity analysis of production and allocation*. Wiley, New York, pp 374–376
- Camerer C (2008) Behavioral game theory. In: Blume LE, Durlauf SN (eds) *The new Palgrave dictionary of economics*, 2nd edn. Mcmillan, New York
- Cournot A (1838) *Researches into the mathematical principles of the theory of wealth*. Kelley, New York. Translation from the French by Nathaniel T. Bacon. Translation publication date: 1960
- Dawid AP (1985) The impossibility of inductive inference. *J Am Stat Assoc* 80(390):340–341
- Dekel E, Feinberg Y (2006) Non-Bayesian testing of an expert. *Rev Econ Stud* 73:893–906
- Feinberg Y, Stewart C (2008) Testing multiple forecasters. *Econometrica* 76(3):561–582
- Foster D, Vohra R (1997) Calibrated learning and correlated equilibrium. *Games Econ Behav* 21:40–55
- Foster D, Vohra R (1998) Asymptotic calibration. *Biometrika* 85:379–390
- Foster D, Young P (2001) On the impossibility of predicting the behavior of rational agents. *Proc Natl Acad Sci* 98:12848–12853

- Foster D, Young P (2003) Learning, hypothesis testing, and Nash equilibrium. *Games Econ Behav* 45:73–96
- Freedman D (1965) On the asymptotic behavior of Bayes estimates in the discrete case II. *Ann Math Stat* 36:454–456
- Fudenberg D, Kreps D (1993) Learning mixed equilibria. *Games Econ Behav* 5(3):320–367
- Fudenberg D, Levine DK (1993a) Self-confirming equilibrium. *Econometrica* 61(3):523–545
- Fudenberg D, Levine DK (1993b) Steady state learning and Nash equilibrium. *Econometrica* 61(3):547–574
- Fudenberg D, Levine DK (1995) Universal consistency and cautious fictitious play. *J Econ Dyn Control* 19:1065–1089
- Fudenberg D, Levine DK (1998) *Theory of learning in games*. MIT Press, Cambridge, MA
- Fudenberg D, Levine DK (1999) Conditional universal consistency. *Games Econ Behav* 29:104–130
- Fudenberg D, Levine DK (2006) Superstition and rational learning. *Am Econ Rev* 96:630–651
- Fudenberg D, Levine DK (2009) *Learning and equilibrium*. *Ann Rev Econ* 1:385–420. Harvard University, Cambridge, MA
- Fudenberg D, Takahashi S (2011) Heterogeneous beliefs and local information in stochastic fictitious play. *Games Econ Behav* 71(1):100–120. Harvard University, Cambridge, MA
- Fudenberg D, Tirole J (1991) *Game theory*. MIT Press, Cambridge, MA
- Germano F, Lugosi G (2007) Global Nash convergence of Foster and Young's regret testing. *Games Econ Behav* 60:135–154
- Goeree JK, Holt CA (2001) Ten little treasures of game theory and ten intuitive contradictions. *Am Econ Rev* 91:1402–1422
- Hahn F (1977) Exercises in conjectural equilibrium. *Scand J Econ* 79:210–226
- Harsanyi J (1973) Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points. *Int J Game Theory* 2:1–23
- Hart S, Mas-Colell A (2000) A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68(5):1127–1150
- Hart S, Mas-Colell A (2001) A general class of adaptive strategies. *J Econ Theory* 98:26–54
- Hart S, Mas-Colell A (2003) Uncoupled dynamics do not lead to Nash equilibrium. *Am Econ Rev* 93:1830–1836
- Hart S, Mas-Colell A (2006) Stochastic uncoupled dynamics and Nash equilibrium. *Games Econ Behav* 57:286–303
- Hofbauer J, Sandholm W (2002) On the global convergence of stochastic fictitious play. *Econometrica* 70(6):2265–2294
- Hurkens S (1995) Learning by forgetful players. *Games Econ Behav* 11:304–329
- Jackson M, Kalai E (1999) False reputation in a society of players. *J Econ Theory* 88(1):40–59
- Jordan JS (1991) Bayesian learning in normal form games. *Games Econ Behav* 3:60–81
- Jordan JS (1993) Three problems in learning mixed-strategy Nash equilibria. *Games Econ Behav* 5(3):368–386
- Jordan JS (1995) Bayesian learning in repeated games. *Games Econ Behav* 9:8–20
- Kalai E, Lehrer E (1993a) Rational learning leads to Nash equilibrium. *Econometrica* 61(5):1019–1045
- Kalai E, Lehrer E (1993b) Subjective equilibrium in repeated games. *Econometrica* 61(5):1231–1240
- Kalai E, Samet D (1984) Persistent equilibria in strategy form games. *Int J Game Theory* 13:129–144
- Kalai E, Lehrer E, Smorodinsky R (1999) Calibrated forecasting and merging. *Games Econ Behav* 18(1/2):151–169
- Kets W, Voorneveld M (2008) Learning to be prepared. *Int J Game Theory* 37(3):333–352
- Kuhn HW (1964) Extensive games and the problem of information. In: Dresher M, Shapley LS, Tucker AW (eds) *Contributions to the theory of games*, vol II. *Annals of mathematics studies*, vol 28. Princeton University Press, Princeton, pp 193–216
- Lehrer E (2001) Any inspection is manipulable. *Econometrica* 69(5):1333–1347
- McKelvey R, Palfrey T (1995) Quantal response equilibrium for normal form games. *Games Econ Behav* 10:6–38
- Milgrom P, Roberts J (1991) Adaptive and sophisticated learning in repeated normal form games. *Games Econ Behav* 3:82–100
- Milgrom P, Weber R (1985) Distributional strategies for games with incomplete information. *Math Oper Res* 10:619–632
- Moulin H (1984) Dominance-solvability and Cournot stability. *Math Soc Sci* 7(1):83–102
- Nachbar JH (2001) Bayesian learning in repeated games of incomplete information. *Soc Choice Welf* 18(2):303–326
- Nachbar JH (2005) Beliefs in repeated games. *Econometrica* 73:459–480
- Nachbar J (2008) Learning and evolution in games: belief learning. In: Blume LE, Durlauf S (eds) *The new Palgrave dictionary of economics*, 2nd edn. Palgrave MacMillan Ltd, New York
- Noguchi Y (2015a) Bayesian learning, smooth approximate optimal behavior, and convergence to ε -Nash equilibrium. *Econometrica* 83(1):353–373. Kanto Gakuin University, Tokyo
- Noguchi Y (2015b) Merging with a set of probability measures: a characterization. *Theor Econ* 10(2):411–444. Kanto Gakuin University, Tokyo
- Nyarko Y (1994) Bayesian learning leads to correlated equilibria in normal form games. *Economic Theory* 4:821–841
- Nyarko Y (1998) Bayesian learning and convergence to Nash equilibria without common priors. *Economic Theory* 11(3):643–655
- Oakes D (1985) Self-calibrating priors do not exist. *J Am Stat Assoc* 80(390):339
- Sanchirico CW (1996) A probabilistic model of learning in games. *Econometrica* 64(6):1375–1393

- Sandholm W (2007a) Evolution in Bayesian games II: stability of purified equilibrium. *J Econ Theory* 136:641–667
- Sandholm W (2007b) Evolutionary game theory. In: *Encyclopedia of complexity and systems science*. Springer. Forthcoming
- Sandroni A (1998) Necessary and sufficient conditions for convergence to Nash equilibrium: the almost absolute continuity hypothesis. *Games Econ Behav* 22:121–147
- Sandroni A (2000) Reciprocity and cooperation in repeated coordination games: the principled-player approach. *Games Econ Behav* 32(2):157–182
- Sandroni A (2003) The reproducible properties of correct forecasts. *Int J Game Theory* 32(1):151–159
- Shapley L (1962) Some topics in two-person games. *Ann Math Stat* 5:1–28
- Voorneveld M (2004) Preparation. *Games Econ Behav* 48:403–414
- Voorneveld M (2005) Persistent retracts and preparation. *Games Econ Behav* 51:228–232
- Young P (1993) The evolution of conventions. *Econometrica* 61(1):57–84
- Young P (2008a) Adaptive heuristics. In: Blume LE, Durlauf S (eds) *The new Palgrave dictionary of economics*, 2nd edn. Palgrave MacMillan Ltd
- Young P (2008b) Stochastic adaptive dynamics. In: Blume LE, Durlauf S (eds) *The new Palgrave dictionary of economics*, 2nd edn. Palgrave MacMillan Ltd.

Fair Division

Steven J. Brams¹ and Christian Klamler²

¹Department of Politics, New York University,
New York, NY, USA

²Institute of Public Economics, University of
Graz, Graz, Austria

Article Outline

Glossary
Introduction
Cutting Cakes
Conclusion
Future Directions
Bibliography

Glossary

Efficiency An allocation is efficient if there is no other allocation that is strictly preferred by at least one agent and not worse for any other agent.

Envy-freeness An allocation is envy-free if no agent strictly prefers any of the other agents' portions.

Equitability An allocation is equitable if every agent values his or her portion the same.

Maximality An allocation is maximin if the worst ranked item received by any of the agents is as highly ranked as possible.

Proportionality An allocation is proportional if each of the n agents gets at least $1/n$ of the good or goods in his or her valuation.

Introduction

Over the last decades, fairness has become a major issue in many different research areas such as economics and computer science. Various books (e.g., Brams and Taylor 1996; Robertson and Webb 1998; Moulin 2003) and surveys (e.g., Thomson 2016; Bouveret et al. 2016; Procaccia

2016) have given many results. The goal in this survey is to single out certain issues of fairness and discuss some contributions. In particular, we focus on (i) cake-cutting, i.e., the division of a single heterogeneous divisible good, and (ii) the allocation of indivisible items.

Issues of fairness come up in many different situations, be it a divorce settlement, the division of land, the sharing of a common resource, the allocation of costs, or simply the division of a birthday cake. From a philosophical point of view, fairness concepts have been widely studied (see, e.g., Rawls 1971; Roemer 1996; Ryan 2006; Young 1994). But how can we transform this work into concepts useful in disciplines such as economics or computer science? The goal of this review is to address certain specific situations in which fairness considerations might play a role and how normative and/or algorithmic concepts can help in finding acceptable fair-division rules and/or solutions.

The above situations and most other fair-division problems do have certain aspects in common. First, there are at least two agents. These might be human beings, but they could also be states, firms, or other entities. Second, there is a resource that is going to be divided. Such a resource can be a heterogeneous good (such as a cake with different toppings), a set of indivisible items, costs or benefits of a joint project, or anything else that might have to be divided or allocated in a division process. Third, there is information about the agents' preferences over the object(s) to be divided. Such preferences could be in the form of value functions (as in cake-cutting), an ordinal ranking over sets of objects (such as in the division of indivisible items), or simply the idea that a smaller cost share is preferred to a larger cost share. Many other aspects of an allocation problem might, of course, be of importance in different situations. For an overview, see Thomson (2016).

In general, the fairness of a division procedure is determined by the properties it satisfies, i.e., a normative approach is followed. Different fair-

division situations might ask for different such properties. However, because the existence of a fair-division procedure or solution does not by itself guarantee that one can also find such a procedure or solution, an important part of the literature is concerned with the algorithmic aspects of fair-division procedures. In this case, an additional problem arises with respect to the computational complexity of fair-division procedures, i.e., even if we have a procedure at hand to deal with a certain problem, will it still be able to find a solution in reasonable time when the fair-division problem increases in size (i.e., the number of agents and/or the resource grows)? Recently, various book-length studies have been published dealing with computational aspects in Economics (see Rothe 2016) and Social Choice Theory (see Brandt et al. 2016).

Cutting Cakes

A cake is usually taken as a metaphor for a single heterogeneous good. How to divide such a good dates back about 2,800 years, when in Hesiod's *Theogony* the Greek gods Prometheus and Zeus argued about how to divide an ox. Eventually they agreed on the following division procedure: Prometheus divided the ox into two piles, and Zeus chose one (see Brams and Taylor 1996). Their procedure, now called **cut-and-choose**, can be seen as the standard example for two-agent situations in the fair-division literature. Many other situations can be analyzed in such a framework, such as the division of land or time slots for the use of a machine, to mention just a few.

First attempts to tackle the cake-cutting problem came from Polish mathematicians in the 1940s, in particular by Hugo Steinhaus, Stefan Banach, and Bronislaw Knaster (see Brams and Taylor 1996; Brams 2006 for a historical overview). Moreover, many books and surveys have been written about this topic; see, e.g., Brams and Taylor (1996), Robertson and Webb (1998), and Procaccia (2013, 2016).

In principle, the cake is seen as the 0–1 interval that has to be divided by n agents who have different valuations of the respective parts of the

interval. (Formally, the agents have value functions, e.g., represented by a probability distribution function, over the interval.) The goal is to divide the interval into pieces, sometimes requiring them to be contiguous, which would imply the use of the minimal number of cuts, i.e., $n - 1$. In the case that an agent does not receive one connected piece, the usual assumption is that the total value of the collection of subpieces is simply the sum of values of the smaller pieces, i.e., additive preferences are assumed. From a normative perspective, the procedures discussed in this literature focus essentially on four properties:

- **Efficiency:** An allocation is efficient if there is no other allocation that is strictly preferred by at least one agent and not worse for any other agent. (This is in the spirit of Pareto optimality, a concept widely used in Economics.)
- **Proportionality:** An allocation is proportional if each of the n agents gets at least $1/n$ of the total cake in his or her valuation.
- **Envy-Freeness:** An allocation is envy-free if no agent strictly prefers any of the other agents' pieces.
- **Equitability:** An allocation is equitable if every agent values his or her piece the same.

Early theoretical results on cake-cutting were mostly concerned with the existence of certain allocations, e.g., Lyapunov's theorem (1940) and results by Dvoretzky et al. (1951) (see also Barbanel 2005). In particular, Lyapunov's theorem shows that there always exists an allocation in which every agent receives a piece that she values at $1/n$ of the total cake and she also values every other agent's piece at $1/n$. However, an agent's piece could consist of a large collection of subpieces, which, from a practical point of view (obtaining only a pile of crumbs), might make such an allocation unsatisfactory.

Let us now turn to cake-cutting procedures (or algorithms) that help in guaranteeing certain fairness aspects in terms of the properties previously defined. Essentially there are two types of procedures, namely, *discrete algorithms* and those that use *moving knives*. Discrete procedures are, in general, algorithms that use a finite number of

discrete steps to reach the final allocation. In moving-knife procedures, on the other hand, an agent has to make continuous decisions and/or valuations while the knife moves along the cake (see Robertson and Webb 1998). In the cake-cutting literature, moving-knife procedures are considerably more complex with respect to the decisions that have to be made by the agents.

As indicated earlier in the example of Prometheus and Zeus when dividing an ox, there exists a simple two-agent cake-cutting procedure called **cut-and-choose**. The procedure works as follows:

1. One agent (the cutter) cuts the cake in two pieces.
2. The other agent (the chooser) chooses one of the two pieces. The remaining piece goes to the cutter. We illustrate the procedure with the following example:⁹

Example 1 Consider a cake (the 0–1 interval) where one half, i.e., the subinterval $[0, \frac{1}{2}]$, is made of strawberry, and the other half, i.e., the subinterval $[\frac{1}{2}, 1]$, is made of chocolate. (As cakes are assumed to be nonatomic, i.e., one single point on the interval does not have any value, we do not have to be concerned about open versus closed intervals.) Assume that agent *A* likes both flavors equally and so only cares about the size of her piece. Agent *B*, on the other hand, does not like strawberry at all and only seeks pieces containing as much chocolate as possible. If those preferences are not known to the agents and agent *A* is the cutter, where should she cut the cake? Obviously, to maximize the minimum share she might receive, she should cut the cake in such a way that both pieces have equal value to her. This leads to her cutting the cake at point $\frac{1}{2}$, generating the two pieces $[0, \frac{1}{2}]$ (containing only strawberry) and $[\frac{1}{2}, 1]$ (containing only chocolate). Agent *B*, the chooser, will now go for the piece that he considers to be larger. Given his preferences, this will be the piece $[\frac{1}{2}, 1]$. Eventually, *A* will get all of the strawberry and *B* all of the chocolate, and, hence, *A* thinks of having a piece of value $\frac{1}{2}$ and *B* a piece of value 1.

Is this a fair allocation? According to our previously defined properties, we see that the allocation is proportional as each of the $n = 2$ agents

gets a piece of value at least $1/n = 1/2$. In addition, it is also envy-free. (In general, envy-freeness implies proportionality: for two agents, envy-freeness and proportionality are equivalent), because no agent strictly prefers the other agent's piece. (Beware that agent *A* is indifferent between her piece and agent *B*'s piece.) However, the allocation is not equitable, because the values that the agents attach to their pieces are not equal ($1/2 \neq 1$). Finally, the allocation is also efficient, because any other allocation that increases the value of one agent's piece would decrease the value of the other agent's piece. (There are situations, i.e., preferences of the agents, in which an allocation determined by one cut will be Pareto dominated by an allocation determined by, e.g., two or more cuts.) Certainly, *A* could be better off by guessing about *B*'s preferences and making the cut at a different point. However, this is risky, because by misjudging the other agent's preferences, she could get a smaller piece (and, hence, risk-averse agents might not want to do this). What can easily be seen, however, is that the allocation does depend on who is the cutter and who is the chooser. If agent *B* were the cutter, he would, to maximize the minimum-size piece that he could get, cut the cake at point $\frac{3}{4}$, i.e., each of the two pieces contains exactly half of the chocolate. In that case *A* would pick the first piece, i.e., interval $[0, \frac{3}{4}]$, and leave the second piece, i.e., interval $[\frac{3}{4}, 1]$, to agent *B*. In that case, *A* receives a value of $3/4$ and *B* a value of $1/2$. Again, this is efficient, proportional, and envy-free, but it is not equitable.

The original cut-and-choose procedure is a discrete algorithm, because it only requires two decisions from the agents. First, the cutter has to cut the cake where he/she thinks the two pieces have the same value. Second, the chooser has to evaluate the two pieces and choose one of them. However, there is also a moving-knife variant of cut-and-choose, introduced by Dubins and Spanier (1961), in which a knife is moved along the interval from 0 to 1 and the first agent to stop the knife receives the piece from 0 to the stopping point and the other agent the remainder. The fairness properties of this procedure are the same as those of the discrete version of cut-and-choose.

Although the cut-and-choose procedure seems promising, it cannot really be extended to more than two agents without losing its fairness properties. If we were to go from two to three agents, guaranteeing envy-freeness with two cuts (the minimal number) becomes difficult. In the 1960s, Selfridge and Conway, however, independently discovered a discrete procedure for three agents using up to five cuts. Following Brams and Taylor (1996), the **Selfridge-Conway procedure** for three agents, A , B and C , can be stated as follows:

1. Let A cut the cake into three pieces she considers of equal value.
2. B trims (if necessary) the piece he considers to be of highest value such that the trimmed piece equals in value his second most preferred piece.
3. C chooses from the three pieces the one she considers of largest value.
4. B chooses from the remaining two pieces. In case they contain the piece that B previously trimmed, he must choose this one.
5. A receives the remaining piece. If there was no trimming by B in step 2, stop the procedure.
6. Because only B or C could have received the trimmed piece, let the one of them, who did not receive it, divide the remainder, which was created in step 2, into three pieces she or he considers to be of equal value.
7. Now, let the agent who previously received the trimmed piece choose first, then A chooses, and finally the cutter of the remainder receives the rest.

Let us show why this leads to an envy-free allocation. Obviously, after the first round of cutting and choosing, the allocation will be envy-free. Because A cuts the cake in the first stage, and therefore will definitely receive one of the untrimmed pieces in step 5, she will not envy either of the other agents. B , as the second chooser, will also not envy anyone, because he receives a tied-for-largest piece, and C , as the first chooser, will get her most preferred piece. Now, let us determine whether the allocation of the remainder of the trimmed piece could make someone envious.

Assume that C received the trimmed piece (the same argument holds if B received it). Hence, because B divides the remainder, he will again receive a tied-for-largest piece. So his total value will be at least as large as the total value of any other player. C is going to choose first. Therefore, she will not envy any of the other agents. Finally, A , the second chooser, will definitely not envy B , because she chooses before him. But she also will not envy C , because A originally preferred her piece in the first round over the piece C received plus the remainder. Thus, the final allocation is envy-free. However, we needed up to five cuts to achieve this allocation.

Actually, guaranteeing fairness in the division of a cake among three or more agents turns out to be difficult if we want to ensure an upper bound on the number of cuts to be made. This computational aspect of fair-division procedures has recently attracted a lot of attention (see, e.g., Brandt et al. 2016 in their *Handbook of Computational Social Choice*). Most of the computational results rely on a framework introduced by Robertson and Webb (1998) in which, focusing on the decisions to be made by the agents, they distinguish between cuts, i.e., cutting the cake at a specific value, and evaluations of a certain sub-piece of the cake. It is the number of those queries which gives an indication of the complexity of a procedure. Using this approach, it has been shown that proportionality is definitely easier to achieve than envy-freeness (see Woeginger and Sgall 2007; Procaccia 2009). Actually, given n players, there exists a procedure that achieves proportionality and uses at most $n \log n$ queries, whereas for any algorithm that leads to envy-freeness at least n^2 queries are necessary. The least complex proportional procedure, as shown by Edmonds and Pruhs (2006), is an algorithm by Evan and Paz (1984) (Evan and Paz 1984 also designed a randomized protocol that uses an expected number of $O(n)$ cuts.). (See also Brams et al. 2011 on an analysis of the divide-and-conquer algorithm.) Essentially it works as follows (for simplicity assume the number of players to be a power of 2):

1. Each agent makes a mark at a point where he/she values the left side of the cake to be

- equal to the right side, i.e., we have a cut query to do this for each of the n agents.
2. Divide the n agents into two subgroups, namely, the first $\frac{n}{2}$ agents with their marks farthest to the left and the other agents with their marks farthest to the right. Cut the cake between those two sets.
 3. Continue by asking the agents to mark the point at which they value the left side of their subpiece equal to the right side of their subpiece. Again, this requires a cut query for each of the n agents.
 4. Repeat this until there are only two agents for each remaining subpiece. Then use cut-and-choose.

As can easily be seen, the Even-Paz procedure leads to a proportional allocation, because each agent in any round is temporarily assigned a subpiece which he/she values at least $1/2$ of the previous subpiece. This, eventually, leads to each agent receiving a piece she values at least $1/n$. Interestingly, the pieces are also contiguous (in contrast to the Selfridge-Conway allocation). However, envy-freeness is not guaranteed as an agent does not have any influence on the allocation in subgroups it does not belong to. How many queries are needed in this procedure? Because in each round each of the n agents makes a decision and there are $\log n$ rounds, the total number of queries is $n \log n$.

As indicated earlier, envy-freeness is much harder to achieve. However, there are some procedures for three and four agents that have been devised and satisfy envy-freeness. Stromquist (1980) introduced a moving-knife procedure that uses four simultaneously moving knives and assigns contiguous pieces. Barbanel and Brams (2004) achieved such an allocation with only two simultaneously moving knives. Extensions to four agents turn out to be even more difficult. Brams et al. (1997) devised a moving-knife procedure which requires up to 11 cuts; however, the extension of the three-agent procedure by Barbanel and Brams (2004) to four agents requires only up to five cuts. Beyond four agents, Su (1999) provided an ε -approximate algorithm that relies on Sperner's lemma, but it requires convergence to an exact division. Essentially, the

only exact moving-knife procedure, introduced by Brams and Taylor (1995), requires an unbounded number of cuts.

This is, from a computational point of view, unsatisfactory. Stromquist (2008) showed that, assuming contiguous pieces, an unbounded algorithm does not exist. Actually, not even the restriction to simpler preference information (such as uniformly distributed valuations of the cake) helps, as was shown by Kurokawa et al. (2013) (see also Brams et al. 2012a for a discussion of structured valuations). However, recently Aziz and Mackenzie (2016) have found a discrete finite and bounded algorithm for n agents. The bounds are, unfortunately, still very large.

Finally, Brams et al. (2013) show, that if one adds equitability as a desirable property to be satisfied, then together with efficiency and envy-freeness such an allocation does not exist for situations with at least three agents. Actually, this is independent of the number of cuts allowed.

If one slightly extends the cake-cutting framework, it is also possible to talk about pies. Whereas cakes are represented by closed intervals, pies are infinitely divisible heterogeneous and atomless one-dimensional continuums whose endpoints are topologically identified. More intuitively, cakes are represented as lines, whereas pies are given as circles. Obviously, because the minimal number of cuts necessary to allocate pieces of a cake to n agents is $n - 1$, for pies this number is n . Barbanel et al. (2009) show that pie-cutting is also a difficult problem, because for three or more players, there exists a pie and corresponding preferences for which no allocation is envy-free and efficient. However, an allocation that is both equitable and efficient always exists. In addition, in a two-agent-setting, in contrast to cake-cutting, pie-cutting allows for positive results with respect to agents being entitled to pieces of a certain size (see Brams et al. 2008).

Dividing Indivisible Items

The previous section on cake-cutting involved one divisible item. For two agents, it was easily shown that an envy-free allocation exists. Of course, there also exist many situations in which we are concerned with the division of different

items. If each of these items were divisible, we could obviously fall back on the cake-cutting setting by just putting the items next to each other (see Jones 2002). The problem is that in many cases items are not divisible, as, e.g., when dividing a painting. Thus, not even in the simplest setting of two agents can we achieve envy-freeness by dividing one indivisible item. The only envy-free allocation would be to throw away the item, but this is, obviously, an inefficient solution. If, eventually, one agent receives the object, then the other agent will envy the receiving agent.

Recently, this branch in the fair-division literature has attracted some attention, as can be seen in a survey paper by Bouveret et al. (2016) and papers by Bouveret and Lang (2008), Bouveret et al. (2010), and Brams et al. (2012b, 2014, 2015, 2017). If the resource to be divided is a set of items, we have to compare different bundles, i.e., we need to go from single items to bundles of items. Of course, we could just let the agents evaluate each of the bundles or provide a ranking of them. However, for just 15 items, this leads to considering $2^{15} = 32,768$ subsets. (A different way to state preferences would be via compact preference representation. This uses a sort of intermediate language as a proxy for representing the agents' preferences. See Bouveret et al. 2016 for an overview.) Although under certain conditions we might reduce the number of relevant comparisons, the general task still seems highly unrealistic. The approach by Herreiner and Puppe (2002) is based on such a ranking of subsets. Their procedure, called *descending demand procedure*, simultaneously goes down the rankings of the agents until, for the first time, the demands of all the agents can jointly be satisfied. (This is related to the fallback bargaining algorithm by Brams and Kilgour 2001.) Although envy-freeness cannot be guaranteed, the goal is to create balanced allocations based on a maximin idea, i.e., maximize the minimum rank of each agent's share.

In the previous approach of evaluating or ranking all subsets, the task for the agents to state those preferences is, in general, considered to be impractical. Other approaches, however, have been devised. A very simple and attractive

procedure to resolve this fair-division problem, called **Adjusted Winner** (AW), was introduced by Brams and Taylor (1996) (see also their popular book, Brams and Taylor 1999). It is based on assigning values to the single items and assuming an additivity condition (hence, there are no complementarities or substitutabilities between items, i.e., the value of an item for the agent does not change irrespective of which other items the agent receives), which makes the value of any set of items to just be the sum of the values of the items contained in that set. In addition, their procedure relies on one important assumption, namely, that eventually one of the items can be divided (without knowing in advance which item this will be).

The procedure works as follows:

1. Each agent distributes 100 points among the items. Given the point distributions, we (provisionally) assign each item to the agent that gave more points to it. If an item received the same number of points from the agents, we allocate it randomly.
2. Determine the sum of points of the items that an agent received. If it is exactly the same for both agents, stop. If the sums are different, transfer the item with the lowest point-value ratio from the agent with the currently higher total sum to the other agent. Calculate the new total sums. In case it is still higher for the first agent, continue by transferring the item with the next-lowest point-value ratio. If it is the same, stop the procedure. If, after the transfer, the receiving agent has a higher total sum, proceed to the last step.
3. Divide the last transferred item so as to equalize the total sums for the agents.

Let us illustrate AW with the following example.

Example 2 Consider a divorce in which there are two agents, *A* and *B*, and six items: the apartment, custody of the kids, the dog, a music collection, jewelry, and a painting. The value that an agent *i* attaches to any item *o* is given by $v_i(o)$. Those values are given in Table 1, where the last column

Fair Division, Table 1 Divorce example

Object	Agent A	Agent B	Ratio $v_B(o)/v_A(o)$
Apartment	35	5	0.14
Custody	25	40	1.6
Dog	20	25	1.25
Music coll.	10	3	0.3
Jewelry	7	15	2.14
Painting	3	12	4

determines the ratio of the value assigned by player B to the value assigned by player A .

Given the point distributions by the agents, we first assign every item to the agent that values it more, i.e., agent A receives the apartment and the music collection and agent B the other four items. As this leads to a total sum of points of 45 for A and 92 for B , we need to transfer the item with the lowest point-value ratio currently belonging to B to agent A . This is the dog. After transferring this item, A has a total value of 65 points and B a total value of 67 points. Hence, an additional transfer needs to be made. The item with the next-lowest point-value ratio is custody; however, we cannot transfer all of it to A but have to divide it so that we equalize the total sums of the players. We do this by solving the following equation: $35 + 10 + 20 + 25\alpha = 15 + 12 + 40(1 - \alpha)$. This leads to $\alpha = 0.03$, i.e., 3% of the item needs to be transferred to A . Practically, this could mean that there are additional visiting rights that agent B concedes to agent A . Eventually, both agents receive a set of items that each of them values at 65.77.

AW is an attractive procedure. As can be shown, it satisfies the main (fairness) properties of efficiency, proportionality, envy-freeness, and equitability. However, it cannot deal with cases in which the item to be divided is not eventually divisible or agents are not able to attach values to each item. Hence, matters turn out to be more complicated when there is no possibility to divide an item or get precise value information from the agents.

Recently, various procedures have been devised to deal with this sort of situation. The framework in which those procedures work relies, in principle, on the possibility of comparing sets

of items based on the ranking of items. (For an introduction to the literature on ranking sets of objects, refer to Barbera et al. 2004.)

Consider a set of four items, 1, 2, 3, and 4, which are ranked by an agent in the following way: $1 \succ 2 \succ 3 \succ 4$. That is, the agent most prefers item 1, then item 2, and so on. How can we compare subsets of these items if the only available information is the agent's preference ranking of the items? If we assume no synergies between the items, i.e., the items are neither complements nor substitutes, then obviously the set $\{1, 2\}$ should be considered better than the set $\{3, 4\}$, because each item in the first set is preferred to every item in the second set. Also, $\{1, 3\}$ can be seen as being better than $\{2, 4\}$, because each item in the first set can be assigned to an item in the second set to which it is preferred. Of course, this is not the case when we compare $\{2, 3\}$ to $\{1, 4\}$, because there is no different item for both 2 and 3 to which they are preferred in the other set.

Obviously, in this context, equitability becomes a meaningless condition. However, envy-freeness can, to a certain extent, still be a relevant property. Bouveret et al. (2010) define the concept of *necessary envy-freeness*. (See also Brams et al. 2003 for a similar definition under a different name.) An allocation (S_A, S_B) , where S_A is the set of items allocated to agent A and S_B the set of items allocated to agent B , is necessarily envy-free whenever for every item in S_A there is an item in S_B that agent A prefers and vice versa for agent B . (This is also called *assuredly envy-free* in Brams et al. 2001.) Hence, there is one major fairness concept that can be ascertained even in this case. Brams et al. (2014) also discuss another interesting property, called *maximality*, which requires the worst ranked item received by any of the agents to be as highly ranked as possible. This seems plausible in situations in which the worst item does have an essential impact on the quality of the set of items received by an agent. An example might be the choice of teams from a pool of workers, where the performance of the team crucially depends on the worker with lowest quality.

To see how algorithms might work in this framework, let us discuss two procedures which

have recently been devised, the **sequential algorithm** (SA) and the **singles-doubles procedure** (SD) introduced by Brams et al. (2014, 2015). Let us start with the apparently simpler SA. It is based on the agents’ rankings of the items and works as follows (beware that any statements about efficiency and envy-freeness require that the agents receive sets of items of the same size; that is why Brams et al. (2015) assume the set of items to be a multiple of the number of agents):

1. On the first round, descend the ranks of the two or more agents, one rank at a time, stopping at the first rank at which each agent can be given a different item (at or above this rank). This is the *stopping point* for that round; the rank reached is its *depth*, which is the same for each agent. Assign one item to each agent in all possible ways that are at or above this depth (there may be only one). This may give rise to one or more SA allocations.
2. On subsequent rounds, continue the descent, increasing the depth of the stopping point on each round. At each stopping point, assign items not yet allocated in all possible ways until all items are allocated.
3. At the completion of the descent, if SA gives more than one possible allocation, choose one that is efficient (Pareto optimal) and, if possible, envy-free.

Let us illustrate SA with the following example:

Example 3 We use the same items as in the AW example. However, consider just the ordinal rankings of the six items and not their precise valuations. This leads to the following preferences given in Table 2 (starting with the most preferred item on top to the least preferred item at the bottom):

The stopping point in round 1 is depth 1, where agent *A* obtains the apartment and agent *B* custody. At depth 2, we cannot give different items to the agents, because agent *B* has already received custody. Hence, in round 2 we must descend to depth 3, to give the agents different items, namely, the dog to *A* and jewelry to *B*. Finally, in round 3, we descend to depth 4 and assign the music collection to *A* and the painting

to *B*. This leads to a final allocation in which the apartment, the dog, and the music collection are assigned to *A* and custody, jewelry, and the painting are assigned to *B*. Notice that this unique allocation is efficient, envy-free, and maximin according to our previous definitions.

Although in the previous example SA provides a normatively satisfying allocation, this, unfortunately, is not always the case. As shown in Brams et al. (2015), whenever SA provides multiple allocations, certain normative properties cannot be guaranteed for more than one of them. However, it is shown that in a two-agent division problem, SA produces at least one allocation that is efficient and, if an envy-free allocation exists, then SA will give at least one allocation that is envy-free and efficient.

In a similar spirit, Brams et al. (2014) devised an algorithm, called the **singles-doubles procedure** (SD), for a two-agent division problem with an even number of items that restricts the number of possible outcomes and guarantees envy-freeness, if it is possible.

Before stating the algorithm, let us define a couple of concepts. The *maximin rank* *m* is the maximum rank such that every item comes up in either *A*’s or *B*’s ranking. A *single* (for an agent *i*) is an item that only comes up in agent *i*’s ranking before the maximin rank. A *double* comes up in both players’ rankings before the maximin rank. The SD algorithm can now be stated as follows:

1. Determine the maximin rank *m*.
2. Assign to each agent its singles. If all items are allocated, stop the procedure.
3. Identify, for each agent, its most preferred unassigned double. If these are different,

Fair Division, Table 2 Preference rankings

P_A	P_B
Apartment	Custody
Custody	Dog
Dog	Jewelry
Music coll.	Painting
Jewelry	Apartment
Painting	Music coll.

assign them accordingly. If they are the same, identify the agent who can be assigned its second most preferred unassigned double while still satisfying envy-freeness. (This can easily be checked by looking at each rank k in an agent's ranking. An agent is not envious if he/she receives strictly more than half of the items from its ranking up to rank k for any odd $k < n$.) Break any ties at random.

We illustrate the SD algorithm with the following example:

Example 4 Let us again use the same rankings as in the previous example, now presented in Table 3.

The double horizontal line after rank 4 indicates the maximin rank, i.e., every item comes up in either A 's or B 's ranking before the line. First, we want to check for existence of a maximin and envy-free allocation. This can easily be done by checking whether the sets of objects up to every odd rank are the same for both agents or not. If not, a maximin and envy-free allocation exists (for a precise definition of the relevant condition, see Brams et al. 2014). As we see in the preferences of Table 3, the top-ranked items are different for both agents. We further need to check rank 3, where we need to compare agent A 's set of items consisting of the apartment, custody, and the dog with agent B 's set consisting of custody, the dog, and jewelry. Because those sets are different, we proceed with rank 5, which, as can be easily seen, also leads to different sets of items. Hence, an envy-free and maximin allocation exists. Now, let us determine the singles and doubles. As the maximin rank is 4, we know that agent A 's singles must be the

apartment and the music collection, whereas for agent B the singles are jewelry and the painting. Hence, this leaves us with two doubles, namely, custody and the dog. Because both agents prefer custody over the dog, we need to check whether we can assign the dog to one of the two agents without creating envy. This is only possible for agent A . Therefore, we need to assign custody to agent B . The final allocation, therefore, assigns the apartment, the dog, and the music collection to agent A and custody, jewelry, and the painting to agent B . As can easily be checked, this allocation is envy-free and maximin.

Although SA and the SD procedure provide the same outcome for the above rankings, this is, in general, not the case. Moreover, note that SA always outputs a complete allocation, whereas the SD procedure only works whenever an envy-free allocation exists. On the other hand, if an envy-free allocation exists, the SD procedure will only output allocations which are envy-free, whereas SA might also output allocations which do not satisfy this property.

A few other procedures for the two-agents case do exist, e.g., the *undercut procedure* (see also Vetschera and Kilgour 2014 for a discussion of related contested pile methods) by Brams et al. (2012b) (with an extension by Aziz 2015 who simplifies the procedure and allows for more general conditions) or the *Trump rule* by Pruhs and Woeginger (2012). Extending these to more than two agents makes the problem significantly harder. Brams et al. (2001) provide a list of paradoxes showing what can go wrong with respect to efficiency and envy-freeness in determining fair shares for the agents.

Another very simple class of procedures are *picking sequences*. These demand rather little preference information from the agents and require only a sequence according to which the agents choose their items. For example, if we have five items and three agents, A , B , and C , then the sequence $ABCCB$ indicates that agent A chooses an item first, then B chooses an item, then C selects two items, and finally B receives the final item. Computationally, this procedure is very appealing because it only requires small bits of information from the agents. Using

Fair Division, Table 3 Preference rankings

P_A	P_B
Apartment	Custody
Custody	Dog
Dog	Jewelry
Music coll.	Painting
Jewelry	Apartment
Painting	Music coll.

sequences is also helpful in determining whether an allocation is efficient. As shown by Brams and King (2005), an allocation is efficient if and only if it is the product of sincere choices by the agents in some sequence. In case of underlying additive preferences, Bouveret and Lang (2011) have studied sequences with respect to their best social performance in utilitarian and egalitarian settings. Under certain assumptions, it turns out that a sequence of strict alternation (e.g., *ABCABC...* for three agents) maximizes utility for society.

If we turn from envy-freeness to proportionality, other approaches are possible. One is to adapt the cut-and-choose algorithm from cake-cutting (see Budish 2011). Because we are concerned with indivisible items, we need only look at so-called maximin shares, i.e., the highest value that an agent can guarantee for himself or herself in case he/she were to divide the items into n piles and will choose last. Procaccia and Wang (2014) show that there are cases in which such a maximin share will not be achieved. However, they also prove that there is always a guarantee to each agent of receiving at least $2/3$ of the maximin share.

Conclusion

In this survey, we have discussed two important kinds of fair-division, cake-cutting, and the division of indivisible items. Various normative concepts and applicable procedures have been introduced. Clearly, envy-freeness is the major normative property in this respect, but it cannot always be satisfied. We also discussed recent results from the literature on Computational Social Choice that concern the computational complexity of fair-division algorithms.

Future Directions

Patently, fair division is a hard problem, whatever the things being divided are. While some conflicts are ineradicable, the trade-offs that best resolve them are by no means evident. Neither are the best algorithms for solving fair-division problems, or their computational complexity; we expect

continuing progress to be made in these areas. We also expect practical problems of fair division, ranging from the splitting of the marital property in a divorce to determining who gets what in an international dispute, to be more and more solved by algorithms available for online use (see, e.g., <http://www.spliddit.org> or <http://www.nyu.edu/projects/adjustedwinner/>).

Bibliography

- Aziz H (2015) A note on the undercut procedure. *Soc Choice Welf* 45(4):723–728
- Aziz H, Mackenzie S (2016) A discrete and bounded envy-free cake cutting protocol for any number of agents. preprint, <https://arxiv.org/abs/1604.03655>
- Barbanel JB (2005) The geometry of efficient fair division. Cambridge University Press, Cambridge
- Barbanel JB, Brams SJ (2004) Cake division with minimal cuts: envy-free procedures for three persons, four persons, and beyond. *Math Soc Sci* 48(3):251–269
- Barbanel JB, Brams SJ, Stromquist W (2009) Cutting a pie is not a piece of cake. *Am Math Mon* 116(6):496–514
- Barbera S, Bossert W, Pattanaik PK (2004) Ranking sets of objects. In: Barbera S, Hammond PJ, Seidl C (eds) *Handbook of utility theory*, vol 2. Springer, New York
- Bouveret S, Lang J (2008) Efficiency and envy-freeness in fair division of indivisible goods: logical representation and complexity. *J Artif Intell Res* 32:525–564
- Bouveret S, Lang J (2011) A general elicitation-free protocol for allocating indivisible goods. *Proceedings of the 22nd international joint conference on artificial intelligence*. pp 73–78
- Bouveret S, Endriss U, Lang J (2010) Fair division under ordinal preferences: computing envy-free allocations of indivisible goods. *Proceedings of the 19th European conference on artificial intelligence*
- Bouveret S, Chevaleyre Y, Maudet N (2016) Fair allocation of indivisible goods. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) *Handbook of computational social choice*. Cambridge University Press, New York
- Brams SJ (2006) Fair division. In: Weingast BR, Wittman D (eds) *Oxford handbook of political economy*. Oxford University Press, New York
- Brams SJ, Kilgour DM (2001) Fallback bargaining. *Group Decis Negot* 10(4):287–316
- Brams SJ, King DL (2005) Efficient fair division: help the worst off or avoid envy? *Ration Soc* 17(4):387–421
- Brams SJ, Taylor AD (1995) An envy-free cake division protocol. *Am Math Mon* 102(1):9–18
- Brams SJ, Taylor AD (1996) Fair division: from cake-cutting to dispute resolution. Cambridge University Press, New York
- Brams SJ, Taylor AD (1999) The win-win solution: guaranteeing fair shares to everybody. W.W. Norton, New York

- Brams SJ, Taylor AD, Zwicker WS (1997) A moving-knife solution to the four-person envy-free cake division problem. *Proc Am Math Soc* 125(2):547–554
- Brams SJ, Edelman PH, Fishburn PC (2001) Paradoxes of fair division. *J Philos* 98(6):300–314
- Brams SJ, Edelman PH, Fishburn PC (2003) Fair division of indivisible items. *Theor Decis* 55(2):147–180
- Brams SJ, Jones MA, Klamler C (2008) Proportional pie-cutting. *Int J Game Theor* 36:353–367
- Brams SJ, Jones MA, Klamler C (2011) Divide and conquer: a proportional, minimal-envy cake-cutting procedure. *SIAM Rev* 53(2):291–307
- Brams SJ, Feldman M, Morgenstern J, Lai JK, Procaccia A (2012a) On maxsum fair cake divisions. *Proceedings of the 26th AAAI conference on artificial intelligence*. pp 1340–1346
- Brams SJ, Kilgour DM, Klamler C (2012b) The undercut procedure: an algorithm for the envy-free division of indivisible items. *Soc Choice Welf* 39:615–631
- Brams SJ, Jones MA, Klamler C (2013) N-person cake-cutting: there may be no perfect division. *Am Math Mon* 120:35–47
- Brams SJ, Kilgour DM, Klamler C (2014) Two-person fair division of indivisible items: an efficient envy-free algorithm. *Not AMS* 61:130–143
- Brams SJ, Kilgour DM, Klamler C (2015) How to divide things fairly. *Math Mag* 88(5):338–348
- Brams SJ, Kilgour DM, Klamler C (2017) Maximin envy-free division of indivisible items. *Group Decis Negot* 26:115–131
- Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) (2016) *Handbook of computational social choice*. Cambridge University Press, New York
- Budish E (2011) The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. *J Polit Econ* 119(6):1061–1103
- Dubins LE, Spanier EH (1961) How to cut a cake fairly. *Am Math Mon* 68:1–17
- Dvoretzky A, Wald A, Wolfowitz J (1951) Relations among certain ranges of vector measures. *Pac J Math* 1:59–74
- Edmonds J, Pruhs K (2006) Cake cutting really is not a piece of cake. *Proceedings of the 17th annual ACM-SIAM symposium on discrete algorithms*. pp 271–278
- Even S, Paz A (1984) A note on cake cutting. *Discret Appl Math* 7:285–296
- Herreiner DK, Puppe C (2002) A simple procedure for finding equitable allocations of indivisible goods. *Soc Choice Welf* 19:415–430
- Jones MA (2002) Equitable, envy-free and efficient cake cutting for two people and its application to divisible goods. *Math Mag* 75(4):275–283
- Kurokawa D, Lai JK, Procaccia A (2013) How to cut a cake before the party ends. *Proceedings of the 27th AAAI conference on artificial intelligence*. pp 555–561
- Lyapounov A (1940) Sur les fonctions-vecteurs complètement additives. *Bull Acad Sci USSR* 4:465–478
- Moulin H (2003) *Fair division and collective welfare*. MIT Press, Cambridge
- Procaccia AD (2009) Thou shalt covet thy neighbor's cake. *Proceedings of the 21st international joint conference on artificial intelligence*. pp 239–244
- Procaccia AD (2013) Cake cutting: not just a child's play. *Commun ACM* 56(7):78–87
- Procaccia AD (2016) Cake cutting algorithms. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) *Handbook of computational social choice*. Cambridge University Press, New York
- Procaccia AD, Wang J (2014) Fair enough: guaranteeing approximate maximin shares. *Proceedings of the 15th ACM conference on economics and computation*. pp 675–692
- Pruhs K, Woeginger GJ (2012) Divorcing made easy. In: Kranakis E, Krizanc D, Luccio F (eds). *FUN 2012*, LNCS 7288. Springer. pp 305–314
- Rawls J (1971) *A theory of justice*. Harvard University Press, Cambridge
- Robertson J, Webb W (1998) *Cake-cutting algorithms*. A K Peters, Natick
- Roemer JE (1996) *Theories of distributive justice*. Harvard University Press, Cambridge
- Rothe J (ed) (2016) *Economics and computation: an introduction to algorithmic game theory, computational social choice and fair division*. Springer, Berlin
- Ryan A (2006) Fairness and philosophy. *Soc Res* 73:597–606
- Stromquist W (1980) How to cut a cake fairly. *Am Math Mon* 87(8):640–644
- Stromquist W (2008) Envy-free cake divisions cannot be found by finite protocols. *Electron J Comb* 15: R11:1–10
- Su FE (1999) Rental harmony: Sperner's lemma in fair division. *Am Math Mon* 106:930–942
- Thomson W (2016) Introduction to the theory of fair allocation. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) *Handbook of computational social choice*. Cambridge University Press, New York
- Vetschera R, Kilgour DM (2014) Fair division of indivisible items between two players: design parameters for contested pile methods. *Theor Decis* 76(4):547–572
- Woeginger GJ, Sgall J (2007) On the complexity of cake cutting. *Discret Optim* 4(2):213–220
- Young HP (1994) *Equity: in theory and practice*. Princeton University Press, Princeton

Social Choice Theory

Salvador Barberà

MOVE, Universitat Autònoma de Barcelona and
Barcelona GSE, Barcelona, Spain

Article Outline

Glossary

Definition of the Subject

Introduction

Cyclical Patterns and Arrow's Impossibility
Theorem

Sen's Result on the Impossibility of a Paretian
Liberal

Incentives: The Gibbard-Satterthwaite Theorem
Escaping Impossibilities

Voting Rules: A Gallery of Proposals

Broader Horizons

Future Research

Bibliography

Keywords

Aggregation rules · Voting methods · Social
choice functions · Impossibility theorems ·
Arrow's impossibility theorem · Chaos
theorems · Characterizations · Strategy-
proofness · Single peakedness · Liberalism

Glossary

Aggregation Rules These are methods that combine information about the preferences of agents in society and turn them into binary relations, interpreted as “collective preferences,” that may or may not inherit the properties of those attributed to individuals.

Arrow's Impossibility theorem This pioneering result expresses the logical impossibility of aggregating individual transitive preferences into social transitive preferences, when a

society faces more than two alternatives, while respecting the Arrowian conditions of Independence of Irrelevant Alternatives, Non-Dictatorship, Universal Domain, and Pareto.

Chaos theorems Cyclical patterns in social preferences arise in many cases, under a wide variety of aggregation rules. In multi-dimensional settings, where social alternatives can be identified with vectors of characteristics, chaos theorems prove that such cyclical patterns can emerge, even if individual preferences are restricted to be saturated and concave, in almost arbitrary forms.

Characterizations Many aggregation rules, social choice functions, and voting rules have been proposed and used by societies. Results that make explicit the unique sets of properties that characterize single rules, or classes of them, illuminate the advantages and the drawbacks of such proposals.

Impossibility theorems Some important results in social choice theory state that certain combinations of axioms cannot be jointly satisfied by any aggregation rule, or by any social choice function, or by any voting method. These results highlight the existence of unavoidable trade-offs between normatively attractive properties that one might demand from such procedures. They are best interpreted as invitations to explore the possibility of satisfying attractive subsets of such properties through the choice of appropriate rules. It requires truthful revelation to be a dominant strategy for each agent, at each possible social situation. Other requirements regarding incentives are also studied in the literature.

Liberalism This is one of the many axioms that can be imposed on social decision-making models. It requires agents to be endowed with power to choose between those alternatives that only differ in aspects that are of their sole concern.

Single peakedness This is a condition on preference profiles, requiring alternatives to be linearly ordered in such a way that, for each individual, any alternative is worse than all

others that are “between” it and his or hers unique best.

Social choice functions These are methods that combine information about the preferences of agents in society and turn them into a social decision.

Strategy-proofness This is a property regarding the incentives of agents to reveal their true characteristics when participating in a collective decision process.

Voting methods These are collective decision-making procedures based on two definitory elements: the information that agents are allowed to express when filling a ballot, and the rules to be used when determining a winner from the set of votes emitted by society members.

Definition of the Subject

The use of voting methods dates back from ancient times, and many rules have been designed and used by different societies to take into account the opinions of their members. Yet, the systematic study of such methods followed a rather discontinued path. It reached one initial peak in the middle ages, lived a golden age during the Enlightenment, and was activated in the middle of the twentieth century thanks to the works of Black and very specially of Arrow, who set the grounds for a comprehensive view of the subject of social choice and provided new tools of analysis. Modern social choice theory is the result of this extended view. Its distinctive characteristic is the normative study of the values and procedures involved in collective decision-making, through the use of the axiomatic method. Kenneth Arrow, John Harsanyi, Amartya Sen, and Eric Maskin are Nobel laureates with leading contributions to social choice.

Introduction

The use of voting methods dates back from ancient times, and many different rules have been designed by different societies to rule their

assemblies, to pass judgment on disputed issues, to appoint officers, or to choose political and religious leaders. A letter of Pliny the Young to Titus Aristo already contains a rich discussion of strategic voting in the Roman Senate's. And the early Byzantine church already used methods that we now call rules of k names, to arbitrate between the secular and the religious powers. Yet, the systematic study of such methods followed a rather discontinuous path. It reached one early peak in the middle age, starting with Ramon Llull's early proposals of voting systems, whose use and study have persisted to our days. Then followed a wide gap, until the Enlightenment marked a golden age, with authors as influential as Borda and Condorcet, who connected voting with the great issues of political philosophy. Then, another period of lesser activity followed, until, in the middle of the twentieth century, the works of Arrow (1951, 1963) and Black (1948, 1958) marked the beginning of a period of renewed interest in voting, and a definite extension of many of the ideas already advanced in its study to cover an even longer list of topics and concerns. The volume on Classics of Social Choice, edited by Mc Lean and Urken (1995), collects many historical contributions to the subject. The editors' introduction highlights the fact that authors in each of these periods of intense discussion were not only distant in time but also mainly ignored its predecessors.

Social choice theory is the result of this extended view. Its topics of interest include the design of collective decision-making mechanisms, whether through voting or by other means, the analysis of procedures to aggregate preferences, information, or judgments and reflect the concern for different ethical and pragmatic principles, ranging from justice and equity to efficiency and incentives and computational complexity. One distinctive characteristic of social choice theory is the use of axiomatic analysis, and its focus on normative, rather than descriptive aspects. This distinguishes social choice from a large part of political science and of political economy, whose choice of topics and methodology are positive. Methods leading, directly or indirectly, to collective decision-

making, are modeled as functions, and their desirable properties are expressed in the form of axioms that each potential method may or may not satisfy. This allows discussing the ability of each method to satisfy some of these desiderata, to compare alternative procedures, and eventually to characterize each one of them in terms of the properties they can meet. Alternatively, when demands on single methods are incompatible, impossibility results are proven and become a helpful tool to stop fruitless debates and to suggest profitable trade-offs between competing principles.

We concentrate in the text on a few essential results that have marked the development of social choice theory, since Arrow to our days, and add a necessarily incomplete list of references. The surveys and textbooks that are mentioned as additional references should help the interested reader to compensate for possible biases in the present list.

Cyclical Patterns and Arrow’s Impossibility Theorem

Consider a society where three voters, 1, 2, and 3 must decide which of three alternatives x , y , and z to choose. Assume that their preferences are expressed by a ranking of those alternatives, as shown in the following table, where alternatives are listed by order of preference for each of the agents (Table 1).

Clearly, x is preferred to y by a majority of voters, and with the same criterion y is preferred to z , and z is preferred to x : The resulting comparisons between these alternatives do not yield a transitive binary relation, but a cyclical one. This simple example is often referred to as the “paradox of voting.” Already known by Condorcet (1785) in the eighteenth century, it announces a host of problems that generalize

to many different contexts and apply to many rules other than simple majority. The example and related ones show that simple majority cannot be used to aggregate the transitive preferences of individuals over more than two alternatives and always guarantee the transitivity of the aggregate (not even its acyclicity, in case social indifferences are allowed). Hence, simple majority is not an aggregator that delivers objects of the same kind than the inputs it aggregates. Moreover, it is not even a proper method to select best decisions. This is because a best alternative may not exist for cyclical binary relations; hence, majority rule cannot always be used to determine which alternative is socially best, a purpose that it could serve if its outputs were transitive, or at least acyclic. Thus, majority voting, which is an extremely attractive rule to aggregate preferences and to adopt decisions when only two alternatives are at stake, runs into difficulties when it comes to properly order, or simply to choose from, sets of three or more alternatives. In his seminal book *Social Choice and Individual Values* (1951, 1963), Arrow extended this remark well beyond the case of majority voting and showed that the problems we just pointed at will be shared by any voting rule satisfying some conditions that he considered minimal.

Let A be the set of alternatives and R be the set of all possible rational (that is, reflexive, complete, and transitive) preference relations on A . For each $i \in N$, let $R_i \subseteq R$ that we shall interpret as the family of i ’s admissible preferences. A *preference aggregation rule* on $\times_i \in_N R_i \subseteq R^n$ is a map, $F: \times_i \in_N R_i \rightarrow B$, where B denotes the set of all reflexive and complete binary relations on A .

Elements in R^n are called *preference profiles* denoted by $R_N = (R_1, \dots, R_n)$.

A preference aggregation rule F has *universal domain* if all individual preferences are admissible. That is, if for each $i \in N$, R_i is the set of all orderings on A (i.e., if $R_i = R$).

A preference aggregation rule F satisfies *independence of irrelevant alternatives* on $\times_i \in_N R_i \subseteq R^n$ if the social preferences between alternatives x and y depend only on the individual

Social Choice Theory,
Table 1 Agents’ preferences

R_1	R_2	R_3
x	z	y
y	x	z
z	y	x

preferences between x and y . Formally, for two preference profiles $(R_1, \dots, R_n)$ and $(R'_1, \dots, R'_n)$ such that for all individuals i , alternatives x and y have the same order in R_i as in R'_i , alternatives x and y have the same order in $F(R_1, \dots, R_n)$ as in $F(R'_1, \dots, R'_n)$.

A preference aggregation rule F satisfies the *Pareto condition* if when an alternative x is ranked strictly above y by all individual orderings $R_1, \dots, R_n$, then x is ranked strictly above y by $F(R_1, \dots, R_n)$.

A preference aggregation rule is *non-dictatorship* if there is no individual, i whose strict preferences always prevail. That is, there is no $i \in N$ such that for all $(R_1, \dots, R_n) \in \times_{i \in N} R_i$, x ranked strictly above y by R_i implies x ranked strictly above y by $F(R_1, \dots, R_n)$, for all x and y .

A *social welfare function* F on $\times_{i \in N} R_i \subseteq R^n$ is a preference aggregation rule **whose elements in the range are transitive orders** (i.e., $F: \times_{i \in N} R_i \rightarrow R$).

Theorem 1 (Arrow's Impossibility Theorem, 1951–1963) When the number of alternatives is larger than two, no social welfare function can simultaneously satisfy the conditions of universal domain, independence of irrelevant alternatives, Pareto and nondictatorship.

Notice that Arrow starts by formalizing the methods that would solve his aggregation problem as functions then imposes desiderata on them, in terms of the axioms that one may want these functions to satisfy, to conclude that these axioms are mutually incompatible. Arrow's theorem can be proved in many different ways. See, for example, Sen (1970), Mas-Colell et al. (1995), Fishburn (1970), Barberà (1983b), Geanakoplos (2005), Yu (2012).

Other results in social choice that are also obtained by the axiomatic method need not be negative or may involve other classes of functions than the ones Arrow referred to. We shall present these different setups and corresponding results, positive or negative.

At any rate, Arrow's impossibility theorem opened the door to numerous interpretations and qualifications and generated a whole body of

literature that we will discuss below. Before we do, let us propose two other important results that also had a strong impact in the development of social choice theory.

Sen's Result on the Impossibility of a Paretian Liberal

Amartya Sen proposed the following puzzle, regarding the possibility of allowing individuals to freely choose the characteristics of alternatives that are of their sole concern, while respecting the Pareto principle. He started by a colorful example. Two agents, one of them (V) lascivious, the other (D) a prude, must decide whether one of them reads a copy of Lady Chatterley's Lover. The alternatives are that V reads it (v), that D reads it (d) and that none of them does (n). Suppose that D 's ideal is that nobody reads, but thinks that, if someone must, he prefers to do it, rather than letting V to enjoy the book. On his side, V considers a waste that nobody reads the book, prefers to read it himself, but enjoys even more the prospect that the prude is the one to do it. The resulting preference orders would be (Table 2)

Sen proposed that a minimal condition of liberalism should allow agents to decide about those aspects of the decision that are of their exclusive concern. In that case, each agent should be able to choose between reading the book or not to, and the social ranking of alternatives should agree with the result of these free choices. Hence, the social preference would have to rank n over d , because D prefers not to read, and v over n , because V prefers to read. If social preferences were transitive, then v would be declared better than d . But this ranking contradicts the Pareto principle, since v is Pareto dominated by d !

Social Choice Theory,
Table 2 Agents' D and
 V preferences

Prude (D)	Lascivious (V)
n	d
d	v
v	n

The contradiction raised by this example was formally stated and extended in Sen (1970) to general situations, in the following terms. Let us assume that each individual $i \in N$ has a “protected or recognized private sphere” D_i consisting of at least one pair of personal alternatives over which this individual is decisive both ways in the social choice process, i.e. $(x, y) \in D_i$ if and only if $(y, x) \in D_i$ with $x = y$. D_i is called symmetric in this case. And decisiveness means that whenever $(x, y) \in D_i$ and xPy , then xPy and whenever $(y, x) \in D_i$ and yPx , then yPx for society.

A preference aggregation rule satisfies *liberalism* if for each individual i , there is at least one pair of personal alternatives $(x, y) \in A$ such that the individual is decisive both ways in the social choice process. Therefore, $(x, y) \in D_i$ and xPy imply xPy and $(y, x) \in D_i$ and yPx imply yPx for society.

Minimal liberalism is the weaker requirement that the above property should at least hold for two agents.

Finally Sen concentrated attention on aggregation rules whose images are cyclic in order to ensure nonempty social choices. The result then is:

Theorem 2 (Sen’s impossibility of a Paretian liberal, 1970) There is no preference aggregation rule generating acyclic social preferences that satisfies the universal domain condition, weak Pareto efficiency, and minimal liberalism.

This “Paretian liberal” paradox gave rise to many comments and reformulations that we’ll discuss below.

Incentives: The Gibbard-Satterthwaite Theorem

The possibility of strategic behavior on the part of voters has been recognized since ancient times: Plinius the Young already discussed an instance of it in one of his letters, and medieval authors were clearly concerned by the possibility that voters might not vote for the best candidate. The idea that one may cast a “useful vote,” rather than a “sincere” one is in the mind of any participant in a decision process. But Arrow explicitly decided

not to tackle the subject, and in spite of some early and brilliant contributions by Farquharson (1969) and Vickrey (1960), it was only in the 1970s that the topic entered in full force into social choice theory, thanks to the simultaneous discovery by Gibbard (1973) and Satterthwaite (1975) of what is now called the Gibbard-Satterthwaite theorem, and the contemporary work of Prasanta Pattanaik (1976, 1978).

Here is the seminal result in this area, again in the form of an impossibility theorem.

A social choice function on $\times_{i \in N} R_i \subseteq R^n$ is a function $f: \times_{i \in N} R_i \rightarrow A$.

A social choice function f on $\times_{i \in N} R_i$ is *nondictatorial* if there is no individual d such that for all $(R_1, \dots, R_n) \in \times_{i \in N} R_i$, $f(R_1, \dots, R_n)$ is a most preferred alternative for R_d in A .

A social choice function f on $\times_{i \in N} R_i$ is *manipulable* at $R_N \in \times_{i \in N} R_i$ by coalition $C \subseteq N$ if there exists $R'_C \in \times_{i \in C} R_i$ ($R'_i \neq R_i$ for any $i \in C$) such that $f(R'_C, R_{N \setminus C}) P_i f(R_N)$ for all $i \in C$. A social choice function is *group strategy-proof* if it is not manipulable at any R_N by any coalition $C \subseteq N$, and it is (individually) *strategy-proof* if it is not manipulable by any singleton agent.

Theorem 3 (Gibbard-Satterthwaite impossibility result) There is no social choice function that is nondictatorial, strategy-proof and has at least three possible outcomes in the range.

The Gibbard-Satterthwaite theorem shows that when society must eventually choose out of more than two alternatives, using a nondictatorial rule, there will exist preference profiles where an agent would gain from not declaring her true preferences. Telling the truth is not a weakly dominant strategy, because it is not always best. In other terms, conditioning one’s vote on those used by other agents is optimal, and strategic voting will be the rule.

Given its importance, many different proofs of this theorem have been developed. See Gibbard (1973) for the original one, Mas-Colell et al. (1995), Schmeidler and Sonnenschein (1978), Barberà (1983a).

In fact, there exist interesting parallels between Arrow’s and the Gibbard-Satterthwaite theorem,

that have been stressed by many authors, including Satterthwaite (1975), Muller and Satterthwaite (1977, 1985), Reny (2001), Eliaz (2004), and Yu (2013).

Escaping Impossibilities

As we already signaled, the axiomatic method is the distinctive mark of social choice theory. Impossibility results are not its only outputs, although the famous ones we quoted above are very salient. Their importance resides on the fact that they pave the way to possibility results: While setting the limits of what can be achieved, and clearly establishing that no collective decision making rule is by any means perfect, they also indicate the directions in which it is possible to identify some attractive ones. In what follows we describe what are often called “escape routes” away from impossibilities. Some allow for real escapes, others retain a negative flavor, proving the depth of those conflicts that were unveiled by the original results.

Escaping Arrow's Impossibility

Arrow's framework is particular in several ways, as it concentrates on social welfare functions, whose domain consists of profiles of preference orders and whose images must also be transitive. One possible way to escape the impossibility is by changing the domain or the range of definition of the functions we consider. This will be discussed later, but first consider those changes that keep the original framework and only depart from it in some specific aspect.

One of them is to **weaken the requirement of transitivity** that is imposed on the aggregate preferences. After all, transitivity is sufficient for the existence of maximal elements given any finite set of alternatives to choose from, but not necessary. By relaxing the transitivity condition on their image, aggregation rules that satisfy all the rest of Arrow's conditions can be identified and characterized. Mas-Colell and Sonnenschein (1972), Plott (1973), Blair et al. (1976), Sen (1977a), and Blair and Pollack (1979) are early examples of an extensive literature on the subject, still recently

enriched by the work of Bossert and Suzumura (2010).

A second route of escape from the impossibility is found by excluding certain preference profiles from the domain of preference aggregation rules, thus **weakening the universal domain requirement**. The more celebrated criterion to restrict preference profiles is Black's (1948) single peakedness condition.

For any $i \in N$, let $\tau(R_i)$ denote the best alternative of R_i on A , also called its *peak*.

A preference profile R_N is *single peaked* if there exists a linear order $<$ on A such that for each agent $i \in N$ and any two alternatives $x, y \in A$, $\tau(R_i) < x < y$, or $y < x < \tau(R_i)$, then xP_iy .

In many applications, the order relative to which single peakedness is predicated arises naturally from the interpretation of the situation that is modeled. This is the case, for example, when alternatives are political candidates, whose position on a left-right spectrum is agreed upon by all voters, or locations of some public facility on a linear space. In other cases, determining the existence of an order that turns the profile into a single peaked one is itself a question that needs analysis. Hence, one may ask whether there are properties that characterize single peakedness without need to explicitly refer to the underlying order. Indeed there are two, that were discovered with a considerable time gap between them. One condition was identified in seminal papers by Sen (1966) and Sen and Pattanaik (1969), and it involves the ranking of triples of alternatives by triples of agents.

A preference profile R_N satisfies *Condition 1* if for any three agents and each triple of alternatives, there exists one alternative that no agent ever ranks as being worse than the other two.

Condition 1 is necessary for a profile to be candidate to satisfy single peakedness, but not sufficient. Actually, it is one of three conditions that Sen and Pattanaik (1969) collected under the common name of value restriction. Much more recently, Ballester and Haeringer (2011) identified a second necessary condition, this time involving four alternatives, but only two agents at a time.

A preference profile R_N satisfies *Condition 2* if for any two agents i and j , and every four

alternatives x, y, z, w such that $xP_iyP_i z$ and $zP_jyP_j x$, it cannot be that wP_iy and wP_jy .

These two authors prove that Conditions 1 and 2, together, characterize single peaked preference profiles. This result nicely closes a gap in our understanding of single peakedness. The authors also characterized, in a similar manner, other related domain restrictions based on the ordering of alternatives.

An immediate consequence of assuming that all preference profiles in the domain of an aggregation rule are single peaked is that the social preference obtained by simple majority voting is transitive, when the number of agents is odd, and quasi transitive in all cases. Since majority voting satisfies all other conditions demanded by Arrow's impossibility theorem, restricting preferences to be single peaked allows us to escape from the impossibility. Other early conditions were due to Inada (1964, 1969). In fact, many other nice aggregation rules can also be defined in single peaked domains. See Austen-Smith and Banks (1999).

A very classical observation regarding **majority voting** as a social choice function is that **for all single peaked preference profiles**, it selects the median of the peak's distribution whenever this is unique. This median voter result is extremely useful to analyze political and location problems, among other applications. Since single peakedness is such a fruitful condition on domains and leads to nice positive results, it is natural to ask whether some alternative condition on preference profiles may lead to similar conclusions.

Arrow stressed that single peakedness relates the agent's preferences to an underlying one-dimensional ranking that expresses a similarity of their views regarding alternatives, even if they value them differently. That led him to interpret that similarity among agents is at the root of the solution to the aggregation problem. Another, not necessarily contradictory order reflecting similarity of views is at the basis of single crossing, which can be expressed as follows.

A preference profile R_N satisfies *single crossing* if there exist a linear order $>$ on the set of alternatives and a linear order $>'$ on the set of agents such that for all $i, j \in N$ such that $j >' i$,

and for all $x, y \in A$ such that $y > x$, if $yP_i x$ then $yP_j x$.

The implications of single crossing on the design of aggregation and decision rules are quite parallel to those of single peakedness. For an odd set of agents, the preference of the median voter, according to the reference ranking of agents, actually coincides with the majoritarian social preference and is thus transitive. And the top alternative for this median agent is the majority winner. Again, slight qualifications must be added when the number of voters is even. Thus, single crossing performs very well.

Given the similarities of results, one could ask whether single peakedness and single crossing are somewhat part of the same family of domain conditions. Indeed, they are. In Barberà and Moreno (2011), it is proven that, in fact, there is a common root: A weaker condition, called top monotonicity, can be imposed on domains, is implied by both single peakedness and single crossing and still retains the common property that a median voter is well defined and would choose the majoritarian outcome.

Still another related domain restriction is intermediateness (see Grandmont (1978) and Gans and Smart (1996)). **Extensions of single peakedness to multidimensional settings** have been proposed. One of them is to impose similar conditions on a tree, rather than on a line, and still helps define domain restrictions allowing for possibility results. See Demange (1982).

However, the most ambitious and probably most natural extension of single peakedness is the one adopted by most of political science and location theory, whereby agent's preferences are concave and have a single maximal element on a multidimensional space. Very strong negative results, sometimes termed "**chaos**" **theorems**, have proven the persistence and pervasiveness of cycles under rules satisfying Arrow's conditions even in these restricted domains. Major early contributions to this very important literature for political science were made by McKelvey (1976, 1979) and Schofield (1978).

For similar reasons, economists have also studied carefully the consequences on Arrow's theorem of imposing domain restrictions that are

standard in the study of exchange economies. Again impossibilities loom over. See Chichilnisky (1980), Le Breton and Weymark (1996), Baigent (2002), Nicolò (2004), and Schummer (1977).

Although different domain restrictions allow to escape from the dire consequences of Arrow's theorem, the more dramatic proof that these results rely on qualitative and not on quantitative restrictions is provided by **single-profile impossibility results**. These were devised in response to an argument of Samuelson (1967) claiming that Arrow's model, with varying preference profiles, is irrelevant to the problem of maximizing a Bergson-Samuelson-type social welfare function (Bergson 1938), which depends on a given set of ordinal utility functions, that is, a fixed preference profile. Single-profile Arrow theorems establish that bad results (dictatorship, or illogic of social preferences, or, more generally, impossibility of aggregation) can be proved with one fixed preference profile (or set of ordinal utility functions), provided the profile is "diverse" enough. For a recent contribution, with references to previous work, see Feldman and Serrano (2008).

A different way to escape Arrow's impossibility is by **assuming that domains do not contain certain quantitative combinations of preferences**. The literature in that direction and also the related calculations of frequencies of violations of the transitive requirement by majority or other aggregation rules are carefully discussed in Gaertner (2001).

Two other requirements imposed by Arrow can be relaxed and help avoid the implication of dictatorship but do not allow for interesting functions to emerge. One is the assumption that the set of voters is finite. Indeed, one can define nondictatorial aggregation rules satisfying all Arrowian conditions, but these rules are severely restricted on how they can distribute the power among voters. See Kirman and Sonderman (1972).

The Pareto principle can also be removed from the list of requirements, while keeping all others, but the only additional rules one obtains that still meet the rest of Arrow's conditions are those where an antidictator always gets the social ranking that results from completely reversing her preference order, as proven in Wilson (1972).

The **weakening of independence of irrelevant alternatives** is an escape route that has generated much controversy about its meaning and its pertinence as a normative requirement. There are certainly many interesting methods to aggregate profiles of individual transitive preferences into transitive social ones, and the Borda rule (Borda 1781) is a salient one, defined as follows. Each agent is expected to submit her entire (strict) preference order over, say, n alternatives. Then each alternative is assigned points that depend on its rank in the preferences of agents: $n - 1$ every time some voter ranks it first, $n - k$ when it ranks it in k -th position, hence 0 when it ranks it last. These points are added, and the total score of each alternative is used to determine an order between the alternatives. Hence, the Borda rule actually provides us with a social welfare function, which actually satisfies all of Arrow's requirements except for independence of irrelevant alternatives.

To see how this condition is violated, just consider the following two profiles and compute the social order according to the Borda rule.

R_1	R_2	R_3
x	z	y
y	x	z
z	y	x

Here, each alternative scores three points in total so we get that the society is indifferent between the alternatives. Replace now R_2 by R'_2 , and notice that the ordering of y and z is the same as before for all agents.

R_1	R'_2	R_3
x	z	y
y	y	z
z	x	x

Now alternatives x, y, z have scores: 2, 4, 3, respectively. Thus, the social preference is yPz , thus contradicts independence of irrelevant alternatives.

Independence of irrelevant alternatives has often been interpreted as a safeguard against the introduction of cardinal considerations and interpersonal comparisons of utility in the formation of social preferences. Hence, authors have found it useful to discuss its role and meaning in a larger framework.

Until here, we have essentially discussed changes in Arrow's demands that allow for some relaxations of his impossibility result, and seen that some departures are more productive than others. But a more radical escape from Arrow's impossibility results comes from changing the very framework of discussion.

By assuming that the inputs to be aggregated were preference orders, rather than utility functions, Arrow explicitly ruled out the possibility of giving the values of utilities any meaning other than ordinal. And he proposed the condition of independence of irrelevant alternatives as a guarantee that interpersonal comparisons of utility would be ruled out. In an important departure from Arrow's framework, Sen (1977b), d'Asprémont and Gevers (1977), Hammond (1976), and Roberts (1980) proposed the study of social welfare functionals, rather than social welfare functions, as a wider framework to analyze the informational basis of aggregation problems. A very lucid account of their main conclusions is contained in Moulin (1988, Chap. 2).

Denote by u the set of all possible utility functions on A .

A *social welfare functional* $F: u^N \rightarrow R$ is a rule that assigns a rational preference relation $F(u_1, \dots, u_n)$ among the alternatives in A to every possible profile of individual utility functions $u_1(\cdot), \dots, u_n(\cdot)$ defined on A . The strict preference relation derived from $F(u_1, \dots, u_n)$ is $F_p(u_1, \dots, u_n)$.

A social welfare functional may use more or less of the many possible pieces of information contained in the numerical values of a utility profile. It is most useful to think of how much information it does not use, by considering the families of preference profiles that share the same social image. Based on the idea that if these families are obtained through certain changes in utilities, the invariance of their images shows that the difference between these utility changes is not taken into account by the functional, the following definitions become natural.

We say that the social welfare functional $F: u^N \rightarrow R$ is *invariant to common cardinal transformations* if $F(u_1, \dots, u_n) = F(u'_1, \dots, u'_n)$ whenever the profiles

of utility functions $(u_1, \dots, u_n)$ and $(u'_1, \dots, u'_n)$ differ only by a common change of origin and units, that is, whenever there are numbers $\beta > 0$ and α such that $u_i(x) = \beta u'_i(x) + \alpha$ for all i and $x \in A$. If the invariance is only with respect to common changes of origin (i.e., we require $\beta = 1$) or of units (i.e., we require $\alpha = 0$), then we say that $F(\cdot)$ is *invariant to common changes of origin or of units*, respectively.

The social welfare functional $F: u^N \rightarrow R$ does not allow *interpersonal comparisons of utility* if $F(u_1, \dots, u_n) = F(u'_1, \dots, u'_n)$ whenever there are numbers $\beta_i > 0$ and α_i such that $u_i(x) = \beta_i u'_i(x) + \alpha_i$ for all i and $x \in A$. If the invariance is only with respect to independent changes of origin (i.e., we require $\beta_i = 1$ for all i) or only with respect to independent changes of units (i.e., we require $\alpha_i = 0$ for all i), then we say that $F(\cdot)$ is *invariant to independent changes of origins or of units*, respectively.

Fundamental rules in the history of economic thought, like utilitarianism, or in the theory of justice, like leximin or maximin, have been characterized on the basis of the different degrees of cardinality or comparability that they allow. For nice simple proofs, see Blackorby et al. (1984).

These tie in with classical and deep philosophical proposals advanced by Harsanyi (1953, 1955) and Rawls (1971), among other thinkers.

Allowing for outputs other than orders to be in the range of a rule is also a possibility: For example, lotteries over preferences, rather than a single one, or probabilistic expressions of binary relations, could be the result of an aggregation exercise. See Fishburn (1973, Chap. 18), Barberà and Sonnenschein (1978), and Barberà and Valenciano (1983).

Finally, let us not forget that the basic assumption precipitating Arrow's impossibility is the existence of at least three alternatives to be ranked. Limiting a priori the range of alternatives to be only two is again a way of escape. The classical paper by May (1952) already contains a characterization of majority voting that works well for the case of two alternatives. See also the first part of Fishburn's (1973) book for a number of additional problems regarding binary choices that are also of interest to collective decision-making.

Escaping Sen's Paretian Liberal Paradox

A good summary of the polemics surrounding Sen's negative result is provided in Gaertner (2009, Chap. 4). Its publication was quickly followed by a large set of proposals, most of them centered on qualifying his axiom of minimal liberalism and discussing the exact implications of his result on the assignment of rights to decide among individuals whose interests conflict, while still respecting his basic framework. The large list of references that one can find in Sen (1982, Chap. 14), including important proposals by Gibbard (1974) among others, proves the seminal character of Sen's work and the quick diffusion of his ideas. Later on, the perspective of commentators and critics shifted, to question whether his original formulation was the appropriate one to debate the issue of liberalism and private rights, and other authors proposed the use of alternative models and tools, based on the study of power distribution in game forms. Major references are Pattanaik (1996), Pattanaik and Suzumura (1996), Gaertner (1986, 1993), and Gaertner et al. (1992).

Escaping the Gibbard-Satterthwaite Impossibility

Like in the preceding cases, one way to obtain positive results regarding the possibility to design strategy-proof rules is by changing the set of objects on which agents can express their preferences and/or the outcomes in the range of the functions under study. One early reaction in that direction was to allow for more than one alternative being chosen by society and for these to have preferences on sets of alternatives. Not much is gained by taking this route. See Barberà (1977), Kelly (1977), Duggan and Schwartz (2000), Benoit (2002), Barberà et al. (2001), and Taylor (2005). Some authors have explored the consequences of relaxing manipulability and studying its costs and benefits (Campbell and Kelly 2009, 2010).

Another possibility, yielding much richer theoretical proposals, though hard to use in practice, is to allow for lotteries over alternatives to be the outcome of voting. See Zeckhauser (1973), Gibbard (1977, 1978), Barberà (1979), Barberà et al. (1998), and Dutta et al. (2002).

When it comes to weaken the assumptions of the theorem within the original framework, the

main possibility (other than restricting the range to only take two values) consists in restricting the domain. And, indeed, domain restrictions are very fruitful in avoiding manipulations.

Again, single peakedness comes to rescue. See Blin and Satterthwaite (1976) and Moulin (1980a). This author characterized all the rules in this domain that are strategyproof and unanimous, i.e., such that, whenever all agents agree on what alternative is their best then this is the social choice. Without loss of generality, we can identify our finite set of ordered alternatives with integer numbers. Then, the class of minmax rules is defined as follows.

A social choice function f is a *minmax rule* associated with a set of integers $(a_S)_S \subseteq \mathbb{N}$ if for each preference profile R_N , $f(R_N) = \min_{S \subseteq N} (\max_{i \in S} \{a_S, \tau(R_i)\})$.

Actually, it is easy to see that all of these rules are not only strategy proof, but also satisfy the stronger condition of group strategy proofness. That fact can be explained by the structure that single peakedness imposes on the domain of admissible preference profiles.

There are other one dimensional domain restrictions that also admit strategy proof rules. Interesting cases are those when preferences are single plateaued or single dipped. See Berga (1998), or Peters et al. (1991, 1992) for an analysis of one-dimensional location problems. Other, more abstract domain restrictions were defined by Kalai and Muller (1977), Kalai and Ritz (1980), and Blair and Muller (1983).

In addition, nondictatorial strategy proof rules can also be defined when alternatives are multi-dimensional and satisfy domains restrictions that extend the notion of single peakedness. But their ability to satisfy strategy-proofness depends very much on how the notion of single peakedness is extended. Results are very negative under Euclidean preferences (see Border and Jordan (1983), Barberà and Peleg (1990) and Peremans et al. (1997)). By contrast, possibilities may be obtained when we extend single peakedness by using the L_1 -norm, as introduced in Barberà et al. (1993).

Consider alternatives to be elements of the Cartesian product of K integer intervals $A = \prod_{k=1}^K B_k$,

where $B_k = \{a_k, \dots, b_k\}$ for any $k = 1, \dots, K$ and endow this set with the L_1 -norm. Given $\alpha, \beta \in A$, the *minimal box containing α and β* is defined by

$$MB(\alpha, \beta) = \{\gamma \in A : \|\alpha - \beta\| = \|\alpha - \gamma\| + \|\gamma - \beta\|\}.$$

A preference $R_i \in R$ is (*multidimensional*) *single peaked* on A if it has a unique maximal element $\tau(R_i) \in A$, and for any $\gamma, \beta \in A$, if $\beta \in MB(\gamma, \tau(R_i))$ then $\beta R_i \gamma$. Denote by P the set of (*multidimensional*) *single peaked preferences* on A .

Notice that the distance between any two alternatives α and β is the length of any shortest path between α and β and the extension of single peakedness is based on the idea that one alternative is better than another if it is “closer” to the best.

Functions defined on Cartesian domains of preferences and Cartesian ranges satisfying the above condition will be strategy proof if and only if they can be described as follows.

A *left coalition system* on B_k is a correspondence C that assigns to any $a_k \in B_k$ a collection of coalitions $C(\alpha_k)$ satisfying the following conditions:

- (1) *Coalition Monotonicity*: if $W \in C(\alpha_k)$ and $W \subset W'$, then $W' \in C(\alpha_k)$.
- (2) *Outcome Monotonicity*: if $\beta_k > \alpha_k$ and $W \in C(\alpha_k)$, then $W \in C(\beta_k)$.
- (3) $C(b_k) = 2^N$.

A *family C of left coalition systems* on B is a collection $\{C_k\}_{k=1}^K$ where each C_k is a left coalition system on B_k . Given $a_k \in B_k$, any $W \in C_k(\alpha_k)$ is a left winning coalition at α_k .

Let $C = \{C_k\}_{k=1}^K$ be a family of left coalition systems on B . The (*multidimensional*) *generalized median voter scheme induced by C* is the social choice function $f: P^n \rightarrow B$ such that: for any $R_N \in P^n$ and any $k \in K$,

$$f_k(R_N) = \min\{\alpha_k \in B_k : \{i \in N : \tau_k(R_i) \leq \alpha_k\} \in C_k(\alpha_k)\}.$$

Theorem 4 *Suppose that the range of the social choice function $f: P^n \rightarrow B$ is a subbox of B . Then, f is strategy-proof if and only if it is a (*multidimensional*) *generalized median voter scheme*.*

A special case of the setup in Barberà et al. (1993) when only two values are possible in each dimension was used by Barberà et al. (1991) to study voting rules to elect members of a club. Many other contributions have been made along similar lines: Serizawa (1999), Le Breton and Sen (1999), Weymark (1999), Nehring and Puppe (2007a, b). Unfortunately, if the domain is restricted and cannot contain some of the conceivable combinations of one-dimensional values, then new and challenging difficulties arise, limiting the variety of voting rules that can still be strategy proof. See Barberà et al. (1997, 2005), Reffgen and Svensson (2012), Reffgen (2015).

Moreover, it is no longer the case that such rules are also group strategy proof, as in the one-dimensional case.

The equivalence between these two conditions depends, once more, on specific features of the domains for which voting rules are defined (see Barberà et al. (2010) and Le Breton and Zaporozhets (2009)). In spite of these caveats, we are far from the negative result of Gibbard and Satterthwaite in multidimensional spaces: This contrasts with the pervasive negative implications of multidimensionality for preference aggregation.

Voting Rules: A Gallery of Proposals

The difficulties in designing the “perfect” voting rule have not deterred individuals or societies from proposing and actually using a rich variety of voting procedures. Notice that although any social choice function can be interpreted as a voting rule, not all voting rules are representable as social choice functions. In general terms, defining a voting rule requires to specify the different ways how voters can fill the ballots, and the form in which these will be scrutinized to attain an outcome. In many cases, voting rules can be identified with social choice functions, provided that a one-to-one correspondence can be established between individual preferences and individual ballots, but in some cases this is not natural, and it is better to think, more in general, of voting rules as game forms, where the ballots play the role of

strategies, and the results of their scrutiny are the outcome function. For an example of voting rules that can be represented directly as social choice function, consider the Borda rule that actually provides us with a social welfare function, which satisfying all of Arrow's requirements except for independence of irrelevant alternatives, and is a member of the large family of scoring rules, all of which assign weights to alternatives based on how voters rank them, and that differ on the way that these weights are defined.

Scoring rules, with an additional tie-breaking criterion when several alternatives attain the same maximal score, also define social choice functions that have been elegantly axiomatized through interesting properties they all share. See Young (1975) and also Saari (1995, 2000).

They all fail, however, to comply with Condorcet consistency, the natural requirement that whenever an alternative would defeat all others in pairwise majority contests, and hence a majority winner exists, it should be chosen. Simple majority was characterized in an early paper by May (1952). Extending its basic principles to choose from more than two alternatives requires to complement it in some way when simple majority leads to a top cycle. Two examples of Condorcet consistent voting rules are the Copeland and the Simpson rule: see Moulin (1988, Chap. 9), for a careful study of these rules, and of the conflict between the principles underlying the two important classes of social choice functions based on scoring and Condorcet consistency, already exhibited in Fishburn (1984).

Other rules based on majorities (simple or qualified) can be defined through the use of **trees**, establishing the order of elimination by vote at each node where an alternative confronts others. Examples of different rules include the amendment and the successive procedure. Apesteguia et al. (2014) provide an axiomatization of these two rules. These and other sequential rules were studied in Banks (1985) and Banks and Bordes (1988). See Moulin 1988 for a nice introduction to the subject and Austen-Smith and Banks (2005) for further details. The use of such methods gives rise to the possibility of agenda manipulation that has been studied in Barberà

and Gerber (2017), following the lead of different political scientists, like Shepsle and Weingast (1984). Moulin and Peleg 1982 studied the possibility of implementing rules through a proper assignment of decision power. Other phenomena, like the no-show paradox, also arise when considering classes of rules.

Yet another seminal idea to construct voting rules consists in using some **metric** and trying to approach the voter's preferences by an appropriate distance-minimizing choice of social preferences or decisions. An important method based on this idea is the Kemeny rule. See Gaertner (2006, Chap. 6).

Approval voting is an interesting method proposed by Brams and Fishburn (1978) that is better described in terms of game forms. The ways in which voters can fill ballots, their strategies, are sets of alternatives, to be interpreted as the list of all alternatives they approve of. Then, the way to arrive at one chosen alternative consists in selecting the one that is approved by most voters (with a tiebreaking rule if necessary). The nice properties of approval voting have given rise to additional proposals. For example, one could admit that voters cannot only approve of some candidates but also disapprove of others. See Felsenthal (1989), Alcantud and Laruelle (2014).

A method to select candidates that has been used since ancient times is based on the idea of balancing the decision power of two different groups of agents, by letting one side to approve of a fixed number k of candidates and then allowing the other side to select one out of them. These rules have been studied by Barberà and Coelho (2010, 2017), who called them **rules of k -names**.

Another proposal on how to choose among candidates has been made recently by Balinski and Laraki (2010), who call it **majority judgment**. These authors present it as a radical departure from standard voting theory. It essentially consists of allowing voters to submit numerical assessments of each candidate, from a given scale, and then choose the candidate that obtains the highest median value.

A different proposal (Casella 2005) involves the use of **storable votes**. It is based on the remark

that voters who participate in making a series of decisions would like to have much to say on those issues they strongly care about, even at the expense of losing some decision power on others. Hence, the idea of endowing agents with storable voting rights that they may use in any amount in order to give more or less support to selected decisions thus allows each voter to express the intensity of its interest regarding different issues.

Another route is suggested by the observation that, in many practical cases, it may be acceptable to **let chance play a role** in adopting a final decision. In fact, randomizing is often a way to provide a decision process with a sense of fairness. We may decide by chance who has to accept a dangerous task or who will benefit from public housing for which there is an excess demand. And in many civilizations and periods, drawing by lot who had to serve in public office was a way to fight corruption. These ideas justify the analysis of voting rules whose outcome is a lottery over alternatives, rather than a single alternative, or even lotteries over social preferences. Thanks to the enlargement of possible outcomes that these extensions provide, voting rules that introduce randomness allow for the solution of different aggregation and incentive issues. They also allow to better understand the reasons why difficulties to design satisfactory rules arise in the subset of cases where only deterministic decisions are allowed. We have already introduced references to the use of lotteries in social choice in a previous section.

Broader Horizons

This already long article has not touched upon a number of very important pieces of literature that are deeply related to the ones we have surveyed and would deserve full articles on their own.

An essential one is the study of **committee decisions**, building from the essential result that is usually known as Condorcet's Jury Theorem. The framework is one where voters are essentially in agreement regarding their objectives but differ on the information they hold and hence on the best course of action that must be taken. The paradigmatic case is that of a jury: All jurors would like to

declare a defendant guilty if he was, and to acquit him if he is innocent, but they may differ on the signals they have received regarding that crucial fact. Discovering truth becomes the objective of voting, a view that is deeply rooted in medieval writings, where the objective is to elicit God's will, or in the Enlightenment attempts to define the will of the people. Young (1988) provides a fascinating account of Condorcet's own internal debate regarding the issue and of his seminal result: When the jury faces two options, the signals they receive are independently drawn and the correct signal is more likely to obtain than the wrong one, simple majority is the maximum likelihood estimator of truth. Extensions of this result are numerous. An important line of research has emphasized the consequences of information exchange among voters and the possible consequences of strategic behavior in that context. See Nitzan and Paroush (1985), Austen-Smith and Banks (1996), and Austen-Smith and Feddersen (2006, 2009).

Another important direction of development is that of **judgment aggregation**. Starting from what is known as the doctrinal paradox, this line of work, essentially developed by philosophers, investigates the difficulties that arise when complex decisions must be adopted, and the aggregation of different pieces of information can follow different patterns. Specifically, the original paradox compares the result of letting each agent aggregate different pieces of information through his own logical reasoning, and then vote on an overall decision, versus the possibility of aggregating each piece of partial evidence through partial votes, and then applying the collective logic to draw a final conclusion. Many aspects are investigated, after this starting point, many of which parallel, and eventually generalize, the classical findings of social choice theory. See List and Pettit (2002), Dietrich (2006), List and Polak (2010), and List (2012).

Voting rules are mechanisms that societies employ to make collective decisions, but certainly not the only ones. When the objects we decide upon have private components, and alternatives are complete descriptions of who gets what, one can hardly choose among them by vote. Markets allocate resources in ways that are different than voting, in spite of analogies made by expressions like "the

money vote” or “voting with the feet.” Rules to allocate children to school, to promote the exchange of kidneys for transplant, to share tasks that no one wants, to ration others that are in excess demand, to split the losses in a bankruptcy according to prior claims are all examples of resource allocation mechanisms that can hardly be considered voting methods. In a sense, voting methods constitute a subclass of mechanisms. The **theory of mechanism design** and the more recent and applied field of **market design** are developed along similar lines than the ones we sketched here, and use to a great extent the axiomatic method that characterizes social choice theory since Arrow’s pathbreaking work. Maskin and Sjöström (2002), Roth (2008), Myerson (2008), Mas-Colell et al. (1995) are references to enter that very large literature.

Other areas of research differ from voting theory but still keep the same intellectual flavor. As pointed out by von Neumann and Morgenstern, comparison is already a form of measuring. Aggregating preferences to evaluate the merits of an alternative over another, or arriving at social decisions after such comparisons, are, after all, attempts to measure the elusive notion of social welfare. **Measurement theory**, as developed by Krantz et al. (1971, 1989), and Luce et al. (1990) is a systematic treatment of issues that arise in the context that we surveyed, but also in others. Social scientists try to measure not only welfare, but also fairness (Young (1994)), inequality (Sen (1992)), and many other magnitudes, including poverty, human development, happiness, or freedom of choice.

Assessing the advantages and the limitations of such measurements through the axiomatic method is very close to what we do when characterizing voting rules according to their properties.

A large number of works are now being written from the perspective of computer scientists, and a new area of **computational social choice** has emerged with force. The models are essentially the same that we have surveyed here, but the questions vary. The complexity of calculations that voters should engage in, to fully grasp the consequences of their actions on collective decisions under different rules, can be assessed with several criteria and is one of the main issues at stake. See Brandt et al. (2016).

Future Research

The study of collective decision-making takes many avenues. Some of them are important and have not been mentioned here: Political economy and much of political science adopt a positive approach, while we have concentrated here in the normative study of issues related to voting through the axiomatic approach. Many of the directions already mentioned will continue to be discussed for sure. Integrating the Arrowian approach with that of the Condorcet jury theorem, probably within a general mechanism design approach, is a challenge. Computational social choice will certainly enlarge the range of ethical and practical issues to worry about, by incorporating and debating on the new possibilities that are open by new technologies on how to gather and to process the people’s opinion in ways that were not accessible before, and are not yet fully exploited. Lab and field experiments are now very much used tools of analysis that will certainly be incorporated to study the performance of voting rules and to control for the importance of different variables. As an example, determining the extent of actual cyclical patterns or the ability of agents to manipulate a given rule may help to determine what aspects of rules deserve closest attention. Behavioral economics has for a long time provided alternative models of individual decision making to the classical assumption that agents are endowed with transitive preferences and always choose the one that is best for them among those that are available. This rational individual’s preferences have been the building block of classical models of collective choice. Formulating and solving the old and the new puzzles that may result from enlarging the concept of rationality, or departing from it, are tasks that lie ahead.

Bibliography

- Alcantud JC, Laruelle A (2014) Disapproval voting: a characterization. *Soc Choice Welf* 43:1–10
- Apesteguia J, Ballester MA, Masatlioglu Y (2014) A foundation for strategic agenda voting. *Games Econ Behav* 87:91–99

- Austen-Smith D, Banks JS (1996) Information aggregation, rationality and the Condorcet Jury Theorem. *Am Polit Sci Rev* 90(1):34–45
- Austen-Smith D, Feddersen T (2006) Deliberation, preference uncertainty and voting rules. *Am Polit Sci Rev* 100(2):209–217
- Austen-Smith D, Feddersen T (2009) Information aggregation and communication in committees. *Philos Trans R Soc B* 364(1518):763–769
- Baigent N (2002) Chapter 18, Topological theories of social choice. In: Arrow KJ, Sen AK, Suzumura K (eds) *Handbook of social choice and welfare*, vol 2. Elsevier, Amsterdam, pp 301–334
- Ballester MA, Haeringer G (2011) A characterization of the single-peaked domain. *Soc Choice Welf* 36:305–322
- Banks JS (1985) Sophisticated voting outcomes and agenda control. *Soc Choice Welf* 1:295–306
- Banks JS, Bordes G (1988) Voting games, indifference, and consistent sequential choice rules. *Soc Choice Welf* 5:31–44
- Barberà S (1977) The manipulation of social choice mechanisms that do not leave too much to chance. *Econometrica* 45(7):1573–1588
- Barberà S (1979) Majority and positional is voting in a probabilistic framework. *Rev Econ Stud* 46(2):379–389
- Barberà S (1983a) Strategy-Proofness and pivotal voters: a direct proof of the Gibbard-Satterthwaite Theorem. *Int Econ Rev* 24(2):413–418
- Barberà S (1983b) Pivotal voters: a simple proof of Arrow's Theorem. In: Pattanaik PK, Salles M (eds) *Social choice and welfare*. North-Holland, Amsterdam, pp 31–35
- Barberà S (2001) An introduction to strategy-proof social choice functions. *Soc Choice Welf* 18:619–653
- Barbera S, Coelho D (2010) On the rule of k names. *Games Econom Behav* 70:44–61
- Barbera S, Coelho D (2017) Balancing the power to appoint officers. *Games Econom Behav* 101:189–203
- Barbera S, Gerber A (2017) Sequential voting and agenda manipulation. *Theor Econ* 12(1):211–247
- Barberà S, Peleg B (1990) Strategy-proof voting schemes with continuous preferences. *Soc Choice Welf* 7:31–38
- Barberà S, Sonnenschein H (1978) Preference aggregation with randomized social orderings. *J Econ Theory* 18(2):244–254
- Barberà S, Valenciano F (1983) Collective probabilistic judgements. *Econometrica* 51(4):1033–1046
- Barberà S, Sonnenschein H, Zhou L (1991) Voting by committees. *Econometrica* 59:595–609
- Barberà S, Gul F, Stacchetti E (1993) Generalized median voter schemes and committees. *J Econ Theory* 61:262–289
- Barberà S, Massó J, Neme A (1997) Voting under constraints. *J Econ Theory* 76(2):298–321
- Barberà S, Bogomolnaia A, van der Stel H (1998) Strategy-proof probabilistic rules for expected utility maximizers. *Math Soc Sci* 35(2):89–103
- Barberà S, Dutta B, Sen A (2001) Strategy-proof social choice correspondences. *J Econ Theory* 101(2):374–394
- Barberà S, Massó J, Neme A (2005) Voting by committees under constraints. *J Econ Theory* 122:185–205
- Barberà S, Berga D, Moreno B (2010) Individual versus group strategy- Proofness: when do they coincide? *J Econ Theory* 145(5):1648–1674
- Barberà S, Moreno B (2011) Top monotonicity: a common root for single peakedness, single crossing and the median voter result. *Games Econom Behav* 73(2):345–359
- Benoît J-P (2002) Strategic manipulation in voting games when lotteries and ties are permitted. *J Econ Theory* 102(2):421–436
- Berga D (1998) Strategy-Proofness and single-plateaued preferences. *Math Soc Sci* 35:105–120
- Bergson A (1938) A reformulation of certain aspects of welfare economics. *Q J Econ* 52(2):310–334
- Black D (1948) On the rationale of group decision making. *J Polit Econ* 56:23–34
- Blackorby C, Donaldson D, Weymark JA (1984) Social choice with interpersonal utility comparisons: a diagrammatic introduction. *Int Econ Rev* 25:327–356
- Blair D, Bordes G, Kelly J, Suzumura K (1976) Impossibility theorems without collective rationality. *J Econ Theory* 13(3):361–379
- Blair DH, Muller E (1983) Essential aggregation procedures on restricted domains of preferences. *J Econ Theory* 30(1):34–53
- Blair D, Pollak R (1979) Collective rationality and dictatorship: the scope of the arrow theorem. *J Econ Theory* 21:186–194
- Blin JM, Satterthwaite MA (1976) Strategy-proofness and single peakedness. *Public Choice* 26:51–58
- Border K, Jordan JS (1983) Straightforward elections, unanimity and phantom voters. *Rev Econ Stud* 50:153–170
- Brams SJ, Fishburn PC (1978) Approval voting. *Am Polit Sci Rev* 72(3):831–847
- Brams SJ, Fishburn PC (2002) Voting procedures. In: Arrow KJ, Sen AK, Suzumura K (eds) *Handbook of social choice and welfare*, vol 1. North-Holland, Amsterdam, pp 173–236
- Campbell D, Kelly J (2002a) Impossibility theorems in the Arrowian framework. In: Arrow K, Sen A, Suzumura K (eds) *Handbook of social choice and welfare*, vol 1. Elsevier Science, Amsterdam
- Campbell DE, Kelly JS (2002b) Impossibility theorems in the Arrowian framework. In: Arrow KJ, Sen AK, Suzumura K (eds) *Handbook of social choice and welfare*, vol 1. North-Holland, Amsterdam
- Campbell DE, Kelly JS (2009) Gains from manipulating social choice rules. *Econ Theory* 40(3):349–371
- Campbell DE, Kelly JS (2010) Losses due to manipulation of social choice rules. *Econ Theory* 45(3):453–467
- Casella A (2005) Storable votes. *Games Econom Behav*, special issue: In Honor of Richard D. McKelvey. Elsevier 51(2):391–419

- Chichilnisky G (1980) Social choice and the topology of spaces of preferences. *Adv Math* 37:165–176
- D'Aspremont C, Gevers L (1977) Equity and the informational basis of collective choice. *Rev Econ Stud* 44:199–209
- de Borda JC (1781) Mémoire sur les élections au scrutin. *Hist Acad Roy Sci*:657–665
- de Condorcet M (1785) Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix, Paris
- Demange G (1982) Single-peaked orders on a tree. *Math Soc Sci* 3(4):389–396
- Dietrich F (2006) Judgment aggregation: (im)possibility theorems. *J Econ Theory* 126(1):286–298
- Duggan J, Schwartz T (2000) Strategic manipulability without resoluteness or shared beliefs: Gibbard–Satterthwaite generalized. *Soc Choice Welf* 17(1):85–93
- Dutta B, Peters H, Sen A (2002) Strategy-proof probabilistic mechanisms in economies with pure public goods. *J Econ Theory* 106(2):392–416
- Dutta B, Peters H, Sen A (2007) Strategy-proof cardinal decision schemes. *Soc Choice Welf* 28(1):163–179
- Eliasz K (2004) Social aggregators. *Soc Choice Welf* 22:317–330
- Farquharson R (1969) *Theory of voting*. Yale University Press, New Haven
- Feldman A, Serrano R (2008) Arrow's impossibility theorem: two simple single-profile versions. *Harv Coll Math Rev* 2:46–57
- Felsenthal DS (1989) On combining approval with disapproval voting. *Soc Choice Welf* 34(1):53–60
- Fishburn PC (1970) Arrow's impossibility theorem: concise proof and infinite voters. *J Econ Theory* 2:103–106
- Fishburn PC (1973) *The theory of social choice*. Princeton University Press, Princeton
- Fishburn PC (1984) Discrete mathematics in voting and group choice. *SIAM J Algebra Discrete Methods* 5(2):263–275
- Gaertner W (1986) Pareto, interdependent rights exercising and strategic behaviour. *J Econ, Suppl* 5:79–98
- Gaertner W (1993) Rights and game forms, types of preference orderings and Pareto inefficiency. In: Diewert WE, Spremann K, Stehling E (eds) *Mathematical modelling in economics. Essays in honor of Wolfgang Eichhorn*. Springer, Berlin/Heidelberg/New York
- Gaertner W (2009) *A primer in social choice theory, revised*. Oxford University Press, New York
- Gaertner W, Pattanaik PK, Suzumura K (1992) Individual rights revisited. *Economica* 59:161–177
- Gans JS, Smart M (1996) Majority voting with single-crossing preferences. *J Public Econ* 59:219–237
- Gardenfors P (1973) Positionalist voting functions. *Theor Decis* 4:1–24
- Geanakoplos J (2005) Three brief proofs of Arrow's Impossibility Theorem. *Econ Theory* 26(1):211–215
- Gibbard A (1973) Manipulation of voting schemes: a general result. *Econometrica* 41:587–602
- Gibbard A (1974) A Pareto-consistent libertarian claim. *J Econ Theory* 7(4):388–410
- Gibbard A (1977) Manipulation of schemes that mix voting with chance. *Econometrica* 45:665–681
- Gibbard A (1978) Straightforwardness of game forms with lotteries as outcomes. *Econometrica* 46(3):595–614
- Grandmont JM (1978) Intermediate preferences and the majority rule. *Econometrica* 46:317–330
- Hammond PJ (1976) Equity, Arrow's conditions, and Rawls' difference principle. *Econometrica* 44:793–804
- Harsanyi JC (1953) Cardinal utility in welfare economics and in the theory of risk-taking. *J Polit Econ* 61:434–435
- Harsanyi JC (1955) Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *J Polit Econ* 63:309–321
- Inada KI (1964) A note on the simple majority decision rule. *Econometrica* 32:316–338
- Inada KI (1969) The simple majority decision rule. *Econometrica* 37:490–506
- Inada KI (1970) Majority rule and rationality. *J Econ Theory* 2:27–40
- Kalai E, Muller E (1977) Characterization of domains admitting nondictatorial social welfare functions and nonmanipulable voting procedures. *J Econ Theory* 16(2):457–469
- Kalai E, Ritz Z (1980) Characterization of the private alternatives domains admitting arrow social welfare functions. *J Econ Theory* 22(1):23–36
- Kelly JS (1977) Strategy-proofness and social welfare functions without single-valuedness. *Econometrica* 45(2):439–446
- Kirman AP, Sondermann D (1972) Arrow's theorem, many agents and invisible dictators. *J Econ Theory* 5(2):267–277
- Le Breton M, Sen A (1999) Separable preferences, strategyproofness and decomposability. *Econometrica* 67(3):605–628
- Le Breton M, Weymark J (1996) An introduction to Arrovian social welfare functions on economic and political domains. In: Schofield N (ed) *Collective decision making: social choice and political economy*. Kluwer, Dordrecht
- Le Breton M, Zaporozhets V (2009) On the equivalence of coalitional and individual strategy-proofness properties. *Soc Choice Welf* 33(2):287–309
- List C (2012) The theory of judgment aggregation: an introductory review. *Synthese* 187(1):179–207
- List C, Pettit P (2002) Aggregating sets of judgments: an impossibility result. *Econ Philos* 18(1):89–110
- List C, Polak B (eds) (2010) *Symposium: judgment aggregation*. *J Econ Theory* 145(2):441–638
- Mas-Colell A, Sonnenschein H (1972) General possibility theorems for group decisions. *Rev Econ Stud* 39:185–192
- Maskin E (1995) Majority rule, social welfare functions, and game forms. In: Basu K, Pattanaik PK, Suzumura K (eds) *Choice, welfare and development. Festschrift for Amartya Sen*. Clarendon Press, Oxford
- Maskin E, Sjöström T (2002) Implementation theory. In: Arrow K, Sen AK, Suzumura K (eds) *Handbook of*

- social choice and welfare, vol 1. Elsevier Science, Amsterdam
- May KO (1952) A set of Independent necessary and sufficient conditions for simple majority decision. *Econometrica* 20:680–684
- McKelvey RD (1976) Intransitivities in multidimensional voting models and some implications for agenda control. *J Econ Theory* 12(3):472–482
- McKelvey R (1979) General conditions for global intransitivities in formal voting models. *Econometrica* 47:1085–1111
- McLennan A (1980) Randomized preference aggregation: additivity of power and strategy proofness. *J Econ Theory* 22(1):1–11
- Moulin H (1979) Dominance solvable voting schemes. *Econometrica* 47(6):1337–1351
- Moulin H (1980a) On strategy-proofness and single-peakedness. *Public Choice* 35(4):437–455
- Moulin H (1980b) Implementing efficient, anonymous and neutral social choice functions. *J Math Econ* 7(3):249–269
- Moulin H (1983) The strategy of social choice, Advanced textbooks in economics, vol 18. North-Holland, Amsterdam
- Moulin H (1994) Social choice. In: Aumann RJ, Hart S (eds) *Handbook of game theory with economic applications*, vol 2. Elsevier, North Holland, Amsterdam. pp 1091–1125
- Moulin H, Peleg B (1982) Cores of effectivity functions and implementation theory. *J Math Econ* 10(1):115–145
- Mueller D (1978) Voting by veto. *J Public Econ* 10:57–75
- Muller E, Satterthwaite MA (1977) The equivalence of strong positive association and strategy-proofness. *J Econ Theory* 14:412–418
- Muller E, Satterthwaite MA (1985) Strategy-proofness: the existence of dominant-strategy mechanisms. In: Hurwicz L, Schmeidler D, Sonnenschein H (eds) *Social goals and social organization. Essays in memory of Elisha Pazner*. Cambridge University Press, New York, 131–172
- Myerson RB (2008) Perspectives on mechanism design in economic theory. *Am Econ Rev* 98(3):586–603
- Nehring K, Puppe C (2007a) Efficient and strategy-proof voting rules: a characterization. *Games Econom Behav* 59(1):132–153
- Nehring K, Puppe C (2007b) The structure of strategy-proof social choice: general characterization and possibility results on median spaces. *J Econ Theory* 135(1):269–305
- Nicolò A (2004) Efficiency and truthfulness with Leontief preferences. A note on two-agent, two-good economies. *Rev Econ Des* 8(4):373–382
- Pattanaik PK (1976) Threats, counterthreats and strategic voting. *Econometrica* 44:91–103
- Pattanaik PK (1978) Strategy and group choice. North-Holland, Amsterdam
- Pattanaik PK (1996) On modelling individual rights: some conceptual issues. In: Arrow KJ, Sen AK, Suzumura K (eds) *Social choice reexamined*. Palgrave Macmillan, UK. ISBN: 978-0-312-12741-1
- Pattanaik PK, Peleg B (1986) Distribution of power under stochastic social choice rules. *Econometrica* 54(4):909–921
- Pattanaik PK, Suzumura K (1996) Individual rights and social evaluation. *Oxf Econ Pap* 48:194–212
- Peremans W, Peters H, van der Stel H, Storcken T (1997) Strategy-proofness on Euclidean spaces. *Soc Choice Welf* 14:379–401
- Peters H, van der Stel H, Storcken T (1991) On uncompromisingness and strategy-proofness, Reports in operations research and systems theory-report M 91–15. University of Limburg, Holland
- Peters H, van der Stel H, Storcken T (1992) Pareto optimality, anonymity, and strategy-proofness in location problems. *Int J Game Theory* 21:221–235
- Plott CR (1973) Path independence, rationality and social choice. *Econometrica* 41(6):1075–1091
- Reffgen A (2015) Strategy-proof social choice on multiple and multi-dimensional single peaked domains. *J Econ Theory* 157:349–383
- Reffgen A, And Svensson L-G (2012) Strategy-proof voting for multiple public goods. *Theor Econ* 7(3):663–688
- Reny PJ (2001) Arrow's theorem and the Gibbard-Satterthwaite theorem: a unified approach. *Econ Lett* 70:99–105
- Roberts K (1980) Possibility theorems with interpersonally comparable welfare levels. *Rev Econ Stud* 47:409–420
- Roth A (2008) What have we learned from market design?, Hahn Lecture. *Econ J* 118(527):285–310
- Saari DG (2000) Mathematical structure of voting paradoxes. II positional voting. *Econ Theory* 15(1):55–102
- Samuelson PA (1967) Arrow's Mathematical politics. In: Hook S (ed) *Human values and economic policy*. New York University Press, New York, pp 41–52
- Saporiti A (2009) Strategy-proofness and single crossing. *Theor Econ* 4:127–163
- Satterthwaite MA (1975) Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. *J Econ Theory* 10:187–217
- Satterthwaite MA, Sonnenschein H (1981) Strategy-proof allocation mechanisms at differentiable points. *Rev Econ Stud* 48:587–597
- Schmeidler D, Sonnenschein H (1978) Two proofs of the Gibbard-Satterthwaite theorem on the possibility of a strategy-proof social choice function. In: Gottinger HW, Leinfellner W (eds) *Decision theory and social ethics*. Reidel, Dordrecht, pp 227–234
- Schofield N (1978) Instability in simple dynamic games. *Rev Econ Stud* 45:575–594
- Schummer J (1977) Strategy-proofness versus efficiency on restricted domains of exchange economies. *Soc Choice Welf* 14(1):47–56
- Sen AK (1966) A possibility theorem on majority decisions'. *Econometrica* 34(2):491–499
- Sen AK (1970) The impossibility of a Paretian liberal. *J Polit Econ* 78:152–157
- Sen AK (1977a) Social choice theory: a re-examination. *Econometrica* 45:53–89

- Sen AK (1977b) On weights and measures: informational constraints in social welfare analysis. *Econometrica* 45:1539–1572
- Sen AK, Pattanaik PK (1969) Necessary and sufficient conditions for rational choice under majority decision. *J Econ Theory* 1:178–202
- Serizawa S (1999) Strategy-proof and symmetric social choice functions for public good economies. *Econometrica* 67(1):121–145
- Shepsle KA, Weingast BR (1984) Uncovered sets and sophisticated voting outcomes with implications for agenda institutions. *Am J Polit Sci* 28:49–74
- Slinko A (2002) The asymptotic strategy-proofness of the plurality and the run-off rules. *Soc Choice Welf* 19:313–324
- Smith J (1973) Aggregation of preferences with variable electorate. *Econometrica* 41(6):1027–1041
- Sönmez T (1999) Strategy-proofness and essentially single-valued cores. *Econometrica* 67:677–689
- Taylor AD (2005) *Social choice and the Mathematics of manipulation*. Cambridge University Press, New York
- Vickrey W (1960) Utility, strategy and social decision rules. *Q J Econ* 74:507–535
- Vorsatz M (2007) Approval voting on dichotomous preferences. *Soc Choice Welf* 28(1):127–141
- Weymark JA (1999) Decomposable strategy-proof social choice functions. *Jpn Econ Rev* 50(3):343–355
- Wilson R (1972) Social choice without the Pareto principle. *J Econ Theory* 5:478–486
- Young HP (1974) An axiomatization of Borda's rule. *J Econ Theory* 9:43–52
- Young HP (1975) Social choice scoring functions. *SIAM J Appl Math* 28:824–838
- Young HP (1988) Condorcet's theory of voting. *Am Polit Sci Rev* 82:1231–1244
- Yu NN (2012) A one-shot proof of Arrow's impossibility theorem. *Econ Theory* 50(2):523–525
- Yu NN (2013) A one-shot proof of Arrow's theorem and the Gibbard-Satterthwaite theorem. *Econ Theory Bull* 1(2):145–149
- Zeckhauser R (1973) Voting systems, honest preferences and Pareto optimality. *Am Polit Sci Rev* 67:934–946
- Zhou L (1991) Impossibility of strategy-proof mechanisms in economies with pure public goods. *Rev Econ Stud* 58:107–119
- Balinski M, Laraki R (2010) *Majority judgment: measuring, ranking, and electing*. MIT Press, Cambridge, MA
- Black D (1958) *The theory of committees and elections*. Cambridge University Press, Cambridge
- Bossert W, Suzumura K (2010) *Consistency choice and rationality*. Harvard University Press, Cambridge, MA
- Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) (2016) *Handbook of computational social choice*. Cambridge University Press, Cambridge, UK
- Fishburn P (1973) *The theory of social choice*. Princeton University Press, Princeton
- Gaertner W (2001) *Domain conditions in social choice theory*. Oxford University Press, Oxford, UK.
- Gaertner W (2006) *A primer in social choice theory*. Oxford University Press, Oxford, UK
- Kelly JS (1988) *Social choice theory. An introduction*. Springer, Berlin/Heidelberg/New York
- Krantz DH, Luce RD, Suppes P, Tversky A (1971) *Foundations of measurement, vol I: Additive and polynomial representations*. Academic Press, New York
- Luce RD, Krantz DH, Suppes P, Tversky A (1990) *Foundations of measurement, vol III: representation, axiomatization, and invariance*. Academic, New York
- Mas-Colell A, Whinston MD, Green J (1995) *Microeconomic theory*. Oxford University Press, Oxford (Part 5)
- McLean I, Urken AB (1995) *Classics of social choice*. The University of Michigan Press, Michigan
- Moulin H (1988) *Axioms of cooperative decision making. Econometric society monographs, vol 15*. Cambridge University Press, Cambridge, UK. ISBN: 9780521360555
- Nitzan S, Paroush J (1985) *Collective decision making: an economic outlook*. Cambridge University Press, New York
- Pattanaik PK (1978) *Strategy and group choice*. North-Holland, Amsterdam
- Peleg B (1984) *Game theoretic analysis of voting in committees*. Cambridge University Press, Cambridge
- Rawls J (1971) *A theory of justice*. Harvard University Press, Cambridge, MA. ISBN: 9780674000780
- Saari DG (1995) *Basic geometry of voting*. Springer, Berlin/Heidelberg/New York
- Sen AK (1970) *Collective choice and social welfare*. Holden-Day, San Francisco/Cambridge
- Sen AK (1982) *Choice, welfare and measurement*. Basil Blackwell, Oxford, UK
- Sen AK (1992) *Inequality reexamined*. Harvard University Press, Cambridge, MA
- Suppes P, Krantz DH, Luce RD, Tversky A (1989) *Foundations of measurement, vol II: Geometrical, threshold, and probabilistic representations*. Academic, New York
- Suzumura K (1983) *Rational choice, collective decisions, and social welfare*. Cambridge University Press, Cambridge
- Suzumura K (2016) *Choice, preferences and procedures*. Harvard University Press, Cambridge, MA
- Young P (1994) *Equity: in theory and practice*. Princeton University Press, Princeton

Reference Books

- Arrow KJ (1951, 1963) *Social choice and individual values*, 2nd edn. Wiley, New York
- Arrow KJ, Sen AK, Suzumura K (eds) (2001/2010) *Handbook of social choice and welfare, vols 1 and 2*. North-Holland, Amsterdam
- Austen-Smith D, Banks JS (1999) *Positive political theory I. Collective preference*. The University of Michigan Press, Michigan. ISBN: 9780472068944
- Austen-Smith D, Banks JS (2005) *Positive political theory II: strategy and structure*. University of Michigan Press, Michigan

Voting

Alvaro Sandroni, Jonathan Pogach,
Michela Tincani, Antonio Penta and
Deniz Selman
University of Pennsylvania, Philadelphia,
PA, USA

Article Outline

Definition of the Subject
Introduction
The Collective Choice Problem
Voting Rules
Welfare Economics
Arrow's Impossibility Theorem
Political Ignorance and the Condorcet Jury
Theorem
Gibbard-Satterthwaite Theorem
Political Competition and Strategic Voting
The Common Value Setting with Strategic Agents
Future Directions
Bibliography

Glossary

Arrow's impossibility theorem Arrow's Impossibility Theorem states that there does not exist a complete social ranking over alternatives that meets minimum impositions of egalitarianism and efficiency, No Dictatorship and Pareto Optimality, respectively. Consequently, there is no voting mechanism that can simultaneously satisfy basic notions of egalitarianism and efficiency.

Cost of voting Any sacrifice in utility that voting entails, such as cognitive costs, time cost of going to the polls, etc.

Collective or social choice problem A collective or social choice problem is a setting in which a group of individuals must jointly decide on a single alternative from a set. The outcome of such a problem potentially affects the welfare of all individuals.

Common values A common values problem is one in which agents share preferences over the alternatives, but may differ on their information regarding the attributes of alternatives.

Condorcet jury theorem In a common values setting, the result that under majority voting the correct candidate will (almost always) be elected so long as the following assumptions are satisfied: (1) each voter's belief is correct with a probability higher than half, (2) each voter votes according to her belief and (3) there is a large number of voters. In political science this result is sometimes referred to as *wisdom of crowds*.

Condorcet paradox See voting cycle.

Downsian model of political competition A model in which candidate strategically position themselves on a unidimensional policy space in order to win an election.

Efficiency Efficiency is a broad criterion that may be used to evaluate the social value of alternative outcomes by demanding that society make use of all valuable resources. Examples of efficiency measures are utilitarianism and Pareto Optimality.

Egalitarianism Egalitarianism is a broad criterion used to evaluate the social value of alternative outcomes by demanding that welfare or resources are evenly distributed across the population.

Gibbard-Satterthwaite theorem The Gibbard-Satterthwaite Theorem states that, under some assumptions, every non-dictatorial social choice function is manipulable.

Majority voting Majority voting is a voting rule that stipulates that each agent vote for a single alternative and an alternative that receives more than half of all votes is the collective choice.

Manipulability of a social choice function A social choice function is said to be manipulable if there is some agent who, given the social choice function, prefers to misreport his true preferences. In a voting context, this translates to voting for an alternative different from that which is most preferred.

Median voter theorem The median voter theorem states that if individuals' preferences are single peaked, then there exists an alternative that beats all others in a pairwise majority vote. Single peaked requires that each individual has a bliss point (most preferred alternative) and alternatives are less preferred the further they are from the bliss point. The selected alternative is then the median bliss point and the voter who has this bliss point is the median voter.

No Dictatorship The No Dictatorship criterion of Arrow's desiderata demands that there is no single individual whose preferences always determine those of society. Any egalitarian arrangement must satisfy the no dictatorship criterion, though arrangements that satisfy no dictatorship need not be egalitarian.

Pairwise majority voting rule A pairwise majority voting rule compares each pair of alternatives in a majority vote. Depending on agents preferences and votes, this rule may lead to a voting cycle.

Paradox of voting the puzzle of why there is high voter turnout in large elections, when the probability of any single vote to be determinant for the outcome is extremely small.

Pareto optimality or pareto efficiency An outcome is Pareto Optimal or Pareto Efficient if it is not possible to increase the welfare of one individual without lessening the welfare of another.

Political ignorance A state in which voters are not well informed about the issues and/or candidates they must vote on.

Social choice function A social choice function is a mapping of all individuals' preferences into an alternative, the social choice.

Strategic abstention The tactical decision of an uninformed citizen to abstain in order to allow informed citizens with the same preferences determine the outcome.

Strategic voting Voting by agents who aim to maximize their utility and might do so by misreporting their true preferences over electoral outcomes.

Utilitarianism Utilitarianism is a conception of efficiency that evaluates outcomes on aggregate utility.

Utility of voting Benefit to citizens from voting, usually divided into two components: a **non instrumental** component which includes utility derived from the mere act of voting and not related to the actual outcome of the election, and an **instrumental** component given by the utility of the outcome a voter would induce if determining the outcome, weighted by the probability that his vote determines the outcome.

Voting cycle or Condorcet's paradox A voting cycle or Condorcet's Paradox results when every feasible alternative is beaten by another in a pairwise majority vote. As such, any collective choice is less preferred to some other alternative by more than half of the population.

Voting rule A voting rule is a mapping of votes into a collective choice. Examples of different voting rules include majority voting and plurality voting. The collective choice may vary under alternative voting rules.

Definition of the Subject

Voting is a fundamental mechanism that individuals use to reach an agreement on which one of many alternatives to implement. The individuals might all be affected by the outcome of such a process and might have conflicting preferences and/or information over alternatives. In a voting mechanism, preferences and information are aggregated as individuals submit votes and a voting rule maps the compilation of votes into the alternative that is to be selected. The use of voting as a means of making a group decision dates back at least to ancient Greece, though French Revolutionary contemporaries Condorcet's and Borda's works are among the pioneers of voting theory. Meanwhile, welfare economists such as Bentham suggested formal definitions of socially desirable outcomes. As voting theory and welfare economics evolved, Arrow's result in the middle twentieth century showed that no mechanism, voting or otherwise, can produce outcomes consistent with some of welfare economists' definitions of socially desirable states. Moreover, results in the

early 1970s suggested that voters have incentives to misrepresent their true preferences in elections. Contemporary voting theory has developed new models of strategic behavior to address questions on how political agents behave and which outcomes voting might produce.

Introduction

Voting is one of the most commonly used ways of making collective decisions. Two issues arise when a group of people must find an agreement on a choice that will potentially affect the welfare of all. First, individuals might have conflicting preferences over the set of alternatives they are choosing from. An interesting question is what is the best way to aggregate individual's preferences in order to reach a common decision when preferences conflict. Second, even if agents share the same preferences, they might possess different information about the alternatives. In this situation, an interesting question is what is the best way to aggregate people's conflicting information so as to make the right choice.

The optimal ways to aggregate preferences and information are among the most important questions that the literature on voting has tried to answer. Since voting concerns societies, as opposed to single individuals, this literature relies on works in the field of welfare economics. This branch of economics is primarily concerned with the analysis and the definition of the welfare of a society. There is no agreement among social scientists on a single definition of welfare of a society. However, two concepts are prominent: efficiency, which concerns the minimal waste scarce resources, and egalitarianism, which concerns the equal distribution of those resources. Using these concepts to define the welfare of a society, social scientists have attempted to answer questions on preference and information aggregation in making a collective choice.

Early works on different voting techniques date back to at least the late eighteenth century, with the works of French mathematicians such as Borda (1781) and Condorcet (1776). Their works were among the first to analyze voting procedures

with formal mathematical tools, which have since been prominently used in the literature on voting. Continuing in this tradition, a fundamental result in voting theory is Arrow's Impossibility Theorem. This is a formal treatment of the problem of aggregation of preferences that reached a striking result: there is no way, under some assumptions, to aggregate individuals' preferences that is minimally egalitarian and minimally efficient.

Unlike the problem of aggregation of conflicting preferences, studies on the aggregation of conflicting information obtain positive results. In particular, the Condorcet Jury Theorem finds a mathematical justification for the phenomenon that is known in Political Science as the *wisdom of the crowds*, that is, the observation that democracies seem to be better at making decisions than single individuals. However Condorcet's result relies on the assumption that people vote sincerely and subsequent works on voters' behavior suggest that this is not always the case. A fundamental theoretical result, known as Gibbard-Satterthwaite theorem, shows formally that, under some assumptions, in every election at least one individual has an incentive to vote non sincerely, i.e. to vote strategically.

Gibbard and Satterthwaite's result is a starting point for a branch of the voting literature that deals with strategic voting. Numerous works analyze strategic voting though the use of mathematical tools such as Game Theory. The latter has been used not only to explain voters' behavior, but also to describe competing candidates' behaviors in an election. The research in this game theoretic literature focuses on the outcome of political competition and on the development of a theory of turnout. The former is analyzed through the use of a model of political competition. A main result is that in a two party election, both parties choose the same political platform in order to maximize the probability of winning an election. The latter is motivated by the paradox of voting, which refers to the empirical observation that citizens vote even when the probability that their vote determines the outcome of the election is negligible, such as in large elections. At present, there is no widely accepted theory to explain this phenomenon.

As with the paradox of voting, the voting literature has still many unanswered interesting questions. As will be mentioned in the section on future directions, the voter's behavior is still only partially understood and much needs to be explained. On a broader level, the study of the historical evolution of democracies needs further developments. Much has been written on the subject of voting, and some interesting results have been obtained, but much still needs to be explained.

The Collective Choice Problem

The basic framework for understanding voting is a *collective* or *social choice problem*: a group of agents must reach an agreement on which alternative to select, and this decision potentially affects the welfare of everyone. Such problems could range from "what should be taught in public schools?" to "who should be president?" Consider two possible alternatives, A and B, and a group of individuals who must jointly choose one of the two. It might be the case that some individuals in society prefer A and some prefer B. These conflicting preferences pose a challenge in determining the appropriate social choice. Alternatively, consider a scenario where A and B are different characteristics that two candidates running for a public office may possess. Suppose that all individuals agree that A is more desirable than B, i.e. this a *common values* setting. However, individuals differ in their information; some think that candidate 1 possesses trait A, while others think candidate 2 does. In this case, the challenge is finding the best way to balance the conflicting information in arriving to a collective choice.

There are several ways to resolve the problems of conflicting preferences and information. For example, people can bargain to reach an agreement, or they can fight. A collective choice problem might also be solved through a dictatorship, or even through the toss of a coin. A fundamental mechanism used to resolve these conflicts is an election in which people submit votes on the feasible alternatives and a *voting rule* maps the

collection of votes into a social choice. As can be seen in the next section, there exists a multitude of voting rules.

Before discussing voting rules, one should note the following distinction between elections: those in which the alternatives are policy and those in which the alternatives are candidates. The former is known as *direct democracy*, a common example of which is a *referendum*. In a referendum, citizens vote on a particular proposal, such as the adoption of a new law. The outcome of such an election is then to implement or not implement the proposal. This is in contrast to the case, known as *representative democracy*, where the election is over candidates. The collective choice in such a system is an agent or group of agents with the responsibility of choosing policy. The following voting rules apply to both situations.

Voting Rules

Majority voting is one of the most commonly used voting rules. The rule prescribes that each citizen vote for a single alternative and an alternative becomes the social choice if it receives more than half of all votes. Clearly, when there are more than two alternatives, majority voting does not necessarily produce a social choice.

To ensure a comparison between alternatives, one can resort to *pairwise majority voting rule*, in which alternatives are voted over pair by pair with a majority vote. That is, a majority vote is held between every pair of feasible alternatives and for each majority vote, the winner is deemed socially preferable to the loser. However, this voting rule might generate an intransitive social preference in which society chooses x to y , y to z , but z to x . Consider a three-individual committee comprised of voters 1, 2 and 3 who must choose one of three alternatives, x , y and z . Individual preferences are such that voter 1 prefers x to y to z , voter 2 prefers z to x to y , and voter 3 preferences y to z to x . By pairwise majority voting x beats y , which in turn beats z , which in turn beats x . This intransitivity over alternatives is known as a *voting cycle* or *Condorcet's Paradox*.

In the presence of a voting cycle, pairwise majority voting might not produce an overall winner, as each alternative might be beaten by another. So, an agenda setter could end a cycle by specifying the order in which the alternatives are to be compared in a pairwise vote. However, pairwise majority voting then places all the decision power in the hands of the agenda setter. Returning to the previous example, suppose the agenda prescribes that voting be carried out between alternatives x and y first and the winner is then to be compared with z . In the first round of voting x beats y and z then beats x , so z becomes collective choice. However, the agenda setter could instead choose an initial comparison between y and z . Since y beats z in the first round and x beats y in the second, x would then be the collective choice. Hence, the agenda setter decides the outcome of the election by choosing the order of the pairwise voting.

A *plurality voting rule* is an alternative way to make a collective choice: agents each vote for one alternative and the alternative with the most votes is chosen.

Other voting rules require agents to submit scores or rankings of all available alternatives, rather than just voting for a single one. In a *Borda Count*, agents rank all alternatives assigning the larger numbers to those that are more preferred. The voting rule sums the scores for each alternative across individuals and the alternative with the highest sum is the social choice.

A *supramajority* voting rule stipulates that the ‘status quo’ alternative is chosen unless another alternative receives at least some specified percentage of the vote larger than fifty percent. In the limit, there might be a *unanimity rule* that mandates one hundred percent of the electorate vote for an alternative for it to be chosen against the status quo. Examples of supramajority rules include the passing of constitutional amendments in the United States, where the current constitution is the status quo. Unanimity rules are commonly found in the judicial system in which all jurors must agree on the guilt of the defendant to override the status quo, the presumption of innocence.

Voting rules might also grant veto power to one or more agents. For example, each of the five permanent members of the fifteen member United Nations Security Council has the power to veto resolutions on particular matters. Any collective choice must therefore have the approval of all five permanent members.

The different voting rules are not simply different methods of arriving at the same social choice. Rather, the result of an election depends critically on the voting rule that is used. In fact, an alternative that is the social choice according to one voting rule might be the least preferred under another. For instance, consider alternatives x , y , and z and seven voters, three who prefer x to y to z , two who prefer y to z to x , and two who prefer z to y to x . By pairwise majority voting, x loses to both y and z . In contrast, x beats both y and z in a plurality vote. This suggests that in order to determine which voting rule to use, one must carefully analyze the various resulting outcomes.

To this end, it is useful to identify criteria that allow one to discriminate among the different outcomes produced under various voting rules. There are two main methods of doing this: efficiency and egalitarianism.

Welfare Economics

The analysis of efficiency and egalitarianism of a social state is among the objectives of a discipline called welfare economics. The first criterion for efficiency used by welfare economists dates back at least to Jeremy Bentham (1789) and is known as utilitarianism. According to utilitarianism, the social interest is judged in terms of the total utility of a community. For example, if by moving from arrangement A to arrangement B Mr. 1 benefits more than Ms. 2 suffers, then the movement from A to B is judged as a social welfare improvement. Notice that in order to implement this criterion, the satisfaction intensities of different individuals must be comparable. In the 1930s this criterion was criticized by Lionel Robbins (1938) and other welfare economists who claimed that the comparison of utilities across individuals has no scientific basis. In the 1940s a new criterion was developed

which required no comparison of individual utilities: the Pareto criterion. A social outcome is said to be Pareto Optimal (Pareto Efficient) if there is no other outcome that would benefit at least one individual without hurting anyone else. Consider a scenario where ten dollars must be split among two individuals who value money and there are two alternatives: either person 1 receives five dollars, person 2 four and the remaining dollar is thrown away, or both receive five dollars. Clearly, the first alternative is not Pareto Optimal because person 2 can be made better off and person 1 would remain as well off if 2 is given the dollar that is being thrown away. Notice that the first alternative is also non utilitarian; in fact, the sum total of utilities can not be maximized when valuable resources are thrown away. However, a drawback of this efficiency criterion is that there exist multiple non comparable Pareto Optimal outcomes: any division of the ten dollars among the two individuals is Pareto Optimal as long as no money is thrown away, since to make one person better off one would have to take resources away from the other person. Hence, the Pareto Optimality criterion does not allow one to distinguish among multiple outcomes. Finally, notice that Pareto Optimal outcomes can be extremely non egalitarian: person 1 receiving ten dollars and person 2 zero is an unequal but Pareto Efficient division.

An alternative criterion often used to discriminate among social outcomes is egalitarianism, which focuses on the distribution of welfare across members of a society. One of the abstract principles behind egalitarianism is *the veil of ignorance* (Harsanyi 1953, 1977; Rawls 1971). Consider a situation in which two persons must share a cake. Pareto Optimality does not help in selecting a division: as in the ten dollars example, any division of the cake is Pareto Optimal. Suppose that one of the two persons sharing the cake is asked to cut it in two without knowing a priori which piece she will receive. Her ignorance about which piece she will receive makes her cut the cake in two equal shares, an egalitarian division. Notice that there are a number of ways to define egalitarian outcomes. For example, Rawls's (1971) *maximin rule* suggests that the

social objective should be to maximize the welfare of the worst-off individual. Finally, notice that an egalitarian outcome might be extremely inefficient. Returning to the ten dollar example, an arrangement where person 1 is given nine dollars and person 2 one dollar is not as egalitarian as one where both are given two dollars, though the latter does not make use of more than half of the available resources.

Arrow's Impossibility Theorem

Ideally, one would like to use the normative criteria of social efficiency and egalitarianism to discriminate between different voting rules and select the best one. However a general result known as *Arrow's Impossibility Theorem* (Arrow 1950) shows mathematically that this is impossible. In his seminal work, Arrow shows that there is no voting mechanism that generates a social consensus on the ordering of the different alternatives while satisfying a number of axioms, among which are the weakest forms of efficiency and egalitarianism: Pareto Optimality and No Dictatorship. The latter, which states that no individual always determines preferences of society, is a weak form of egalitarianism. While a non-dictatorial society can be quite unequal, any egalitarian society must be non-dictatorial.

A number of possibility results have been obtained by the relaxation of some of Arrow's axioms. For example, pairwise majority voting with a particular restriction on individual tastes, which violates what Arrow called Unrestricted Domain, satisfies Pareto Optimality and No Dictatorship, while generating an ordering of the social alternatives.

Black (1948) noticed that pairwise majority voting produces an outcome that is not subject to Condorcet's paradox when individual preferences are single-peaked: every individual must have a most preferred alternative (*bliss point*) and between any two alternatives he prefers the one that is closer to his bliss point. An important result in voting theory, called the median voter theorem, shows that when individuals' preferences satisfy this condition, the bliss point of the *median voter*

beats any other alternative by pairwise majority voting. The median voter is found by ordering voters according to their bliss points. The importance of this theorem derives from its ability to describe how democracies work in practice. It is commonly observed that candidates try to appeal to voters who are politically moderate, or “in the middle”: this is consistent with the theory, which suggests that these are the preferences that will eventually prevail in a democratic system.

Political Ignorance and the Condorcet Jury Theorem

As mentioned earlier, voting is not only a way to aggregate conflicting preferences but it is also a way to aggregate individual, possibly conflicting, information when preferences are partially or totally aligned. This is the case in *common value* settings. When people would agree on the best choice if given the same information on alternatives, but differ in the information they actually receive, a natural question is which is the voting mechanism that aggregates information in a way that maximizes the probability of the right decision being made. A result called *Condorcet Jury Theorem* (Condorcet 1976) shows that among all the possible voting rules, simple majority rule guarantees that the right decision is made under three crucial assumptions: that each voter has a correct belief with a probability higher than 50%, the voter votes according to his belief, and that the electorate is very large. Before going into the details of the theorem, it must be mentioned that this result has a practical importance. A number of works document that voters are ignorant over both policy and candidate alternatives over which they are to vote. Campbell et al. (1960) claim that “many people know the existence of few if any of the major issues of policy”, while (Dye and Zeigler 1970) discusses “mass political ignorance” and “mass political apathy” as playing key roles throughout the history of American politics. More recently, the 2004 *American National Election Study* found that Americans performed extremely poorly when asked simple questions about the political system and the leaders in

charge of it. This evidence is in favor of what is called *political ignorance*. In a setting where voters are politically ignorant, the Condorcet Jury Theorem provides a valuable insight as to how much political information matters in determining the outcome.

Consider a committee who has to elect an administrator out of two candidates, one “good” and one “bad.” Assume that all the members of the committee share the same preferences: they all prefer the good administrator to be selected. Individuals differ, however, in the information they have about which candidate is the good one. Suppose that each voter has a belief about which is the good candidate and votes according to his belief. If each voter has a correct belief with probability higher than 50%, then by the Law of Large Numbers as the number of voters becomes very large the probability that more than half of the electorate votes for the right candidate goes to one. Hence, under simple majority the probability that the right choice is made goes to one. Condorcet’s conclusion is that in a common value setting a democratic decision is superior to an individual decision, because each voter makes the wrong decision with a non-negligible positive probability, whereas the population as a whole makes the right decision almost always. As far as political ignorance is concerned, this result shows that it is not necessary for an electorate to be well informed for it to make the right decision. So long as each voter has a correct belief with a probability higher than a half, the electoral outcome will almost always be identical to one in which the electorate was perfectly informed. Therefore, Condorcet’s result implies that the ignorance of individual voters is overcome by the aggregation of information in an election.

Gibbard-Satterthwaite Theorem

Thus far, citizens have been treated as if they disregard any tactical considerations when faced with a voting decision. However, a citizen might find it worthwhile to misrepresent his true preferences in order to achieve a social outcome more preferred than the one that would result if he voted

naively. Consider an election in which a status quo will be replaced if a simple majority agrees on one of three candidates. Suppose the status quo is a conservative government, and the three alternative candidates to be voted for are a conservative, a moderate and a liberal. Imagine that there are only three voters, and two votes have already been cast: one is for the moderate candidate, one is for the conservative one. Suppose that the last individual who is called to vote is politically liberal. He knows that if he votes for his truly most preferred candidate, the liberal one, there would be a tie and the status quo conservative government would not be replaced. However, by voting for his second most preferred alternative, the moderate candidate, he would break a tie and the status quo government would be replaced with a moderate one. A liberal voter prefers this outcome to the one where conservatives win. Therefore he has an incentive to misrepresent his true preferences and vote tactically for his second-best alternative.

A powerful result in voting theory called Gibbard-Satterthwaite theorem (Gibbard 1973; Satterthwaite 1975) shows formally that in most electoral settings at least one citizen has an incentive to vote tactically. Define a *social choice function* as a mapping of all individuals' preferences into a collective choice. The theorem states that there is no social choice function that is Non Dictatorial and Non Manipulable, i.e. such that no agent has an incentive to vote tactically. Consequently, for every voting rule there is at least one agent with an incentive to misrepresent her preferences.

It should be mentioned that the Gibbard-Satterthwaite theorem places also other technical restrictions. Furthermore, in a setting such as a majority vote, a misrepresentation of one's true preferences coincides with voting for an alternative which the voter does not rank top. In the original formulation of their theorem, however, Gibbard and Satterthwaite dealt with mechanisms where agents are required to submit a ranking over all alternatives, and a misrepresentation of tastes in their framework does not coincide necessarily with a misrepresentation of only the most preferred alternative (see (Mas-Colell et al. 1995)).

This result suggests that for a deep understanding of voting one should not focus only on the

assumption that citizens vote sincerely. Tactical voting is not only an abstract possibility, it is also an actual behavior that must be considered in any voting analysis. The following sections explore how the literature on voting has dealt with strategic voting.

Political Competition and Strategic Voting

The early works of Downs (1957) and Tullock (1967) initiated the analysis of political issues within a strategic framework, where voters and/or candidates are assumed to be rational decision makers. *Political competition* describes a situation in which candidates strategically position themselves in order to win an election, whereas *strategic voting* refers to individuals' decision to vote so to maximize utility by sometimes misreporting true preferences. This game theoretical framework developed due to positive arguments such as Gibbard-Satterthwaite theorem, which suggests that voters have an incentive to behave strategically, and is also due to spread of game theory as a dominant tool in economic analysis.

Political Competition

The classic model analyzing the candidates' choice of positioning on the political spectrum is that of Downs (1957), who adapted the classical *Hotelling model* (Hotelling 1929) to the analysis of the choice of political platforms by candidates. In the *Downsian model of political competition*, there is a unidimensional policy space, representing the political spectrum. There are two candidates who position themselves on this policy space. Each voter has a most preferred point on this space and prefers points closer to this point than those further away, i.e. each voter has single-peaked preferences. Downs argues that strategic candidates concerned only with winning position themselves at the point that is most preferred by the median voter. If candidate *A* were positioned anywhere else, say to the left of the median voter, candidate *B* could get the majority of votes by positioning himself between candidate *A* and the median voter; all the agents to the right

of *B*, constituting more than half of the voters, would prefer *B* to *A*. Given this scenario, candidate *A* (for the same reason) would then have an incentive to position himself between *B* and the median voter, and so on. Hence, the result is that both candidates position themselves at the policy platform most preferred by the median voter. An interesting implication of this result is that under a democracy with two parties, both parties act identically, and therefore there are only as many positions (just one) taken by political parties as there would be in a dictatorship. However, it is crucial that two parties exist so that the competition between them can allow the chosen policy point to represent the preferences of the voters. This is in contrast to a dictatorship in which the ruling party can implement its own preferred policy without voter approval.

Notice that the example above assumes a two-party system (for models that allow for more than two parties see (Besley and Coate 1997; Osborne and Slivinski 1996)). (Duverger 1972) posits that in a representative democracy with a plurality voting rule, only two parties compete in the elections. This theory, known as *Duverger's Law*, states that a *proportional representation system*, in which parties gain seats proportional to the number of votes received, fosters elections with numerous parties. In contrast, a plurality system marginalizes smaller parties and results in only two parties entering into political competition.

The Decision to Vote: The Paradox of Voting

Another issue raised by Downs, one which focuses on voters' behavior rather than candidates', is known as *the paradox of voting*. It refers to the fact that in a large election, the probability that any single vote determines the outcome is vanishingly small. If every person only votes for the purpose of influencing the outcome of the election, even a small cost of voting would be sufficient to dissuade anyone from voting. Yet, it is commonly observed that turnout is very high, even in large elections. From the large empirical literature on turnout in elections, some facts seem to be acquired knowledge: (1) turnout is higher in more important elections (e.g., Presidential election in the US have a significantly higher turnout than Gubernatorial elections), (2) turnout is

generally higher in close elections (i.e. with smaller margins of victory), and (3) turnout rates are different among groups with different demographic characteristics. For instance, from the thorough work by (Wolfinger 1980), it emerges that education has a substantial effect on the probability that one will vote. Income has less of an effect once it has been controlled for the impact of other variables. After education, the second most important variable is age, which appears to have a strong positive relationship with turnout. Other socio-economic variables are also important; in particular, racial minorities appear to be less likely to vote. Finally, turnout seems to be significantly influenced by factors such as the weather conditions on the day of the election and voters' distance from the polls (see (Coate and Conlin 2004)).

Such comparative statics suggest that it is appropriate to model voters' behavior as a rational choice problem within a standard utility maximization framework. The modern theory of voting applies the classic utilitarian framework to the voting problem, positing that agents decide whether or not to vote by comparing the *cost of voting* with the *utility of voting*. The traditional starting point for the modern theory of voting is (Riker and Ordeshook 1968), who formalize the insights of (Downs 1957; Tullock 1967) in a simple utilitarian model of voting. The *cost of voting* comprises any sacrifice in utility that voting entails. The *utility of voting* is usually divided into two components: a *noninstrumental* component and an *instrumental* component. The *non-instrumental* component includes utility derived from the mere act of voting and not related to the actual outcome of the election. It may include, for instance, the sense of civic duty. There is considerable evidence that voters are motivated by a sense of civic duty (see, for example, (Blais 2000)). The *instrumental* component is the utility of the outcome a voter induces if her vote determines the outcome, weighted by the probability that her vote actually determines the outcome.

The instrumental component of the utility of voting has attracted most of the attention in the literature. It is typically analyzed through the rational theory of voting, which is motivated by an empirical observation: there exists a strong

positive correlation between turnout rate and closeness of the election. This fact suggests that, *ceteris paribus*, voters are more likely to vote if their vote is more likely to make a difference. The main theoretical problem is to endogenize the probability that each voter is pivotal, i.e. that his vote is determinant for the outcome of the election.

Ledyard (1981, 1984) are among the early works in the literature on game theoretical models of the pivotal-voter. In these models, voters infer the probability of being pivotal from the equilibrium strategies of other voters. Subsequently, they decide whether or not to vote, trading off the cost of voting with the expected (instrumental) utility of voting. Although Ledyard did not focus on the magnitude of turnout in a strategic model, this question is addressed by Palfrey and Rosenthal (1983, 1985) who model elections with uncertainty about the total number of voters. Voters strategically choose whether or not to vote for their favorite alternative amongst two candidates. However, Palfrey and Rosenthal's theories do not explain high turnout in large elections, when the cost of voting is not very (and unrealistically) low.

Ultimately, the game-theoretic approach to costly voting could not escape the *paradox of voting*. Since the probability of being pivotal is very small in large elections, the individual incentives to vote cannot justify high turnouts unless the cost of voting is sufficiently small. Conversely, regardless of how small the cost of voting is, the theory posits that there should be low turnout as the election becomes arbitrarily large, which is in contrast to empirical evidence. The puzzle that remains open is how to reconcile the evidence of high turnout in large elections with the responsiveness of turnout levels to the closeness of the election.

Mobilization and Group-Based Notion of Welfare

Two strands of the literature try to overcome the paradox of voting by focusing on groups of like-minded people rather than on individual agents. These are models of *mobilization* and models incorporating a *group based* notion of welfare.

In models of mobilization, the population of voters is assumed to be divided into groups, each

of which has a leader who has the same preferences as all agents in the group and coordinates their behavior. The turnout decision within each group is determined by how the leaders allocate costly resources to voters. It is as if leaders buy the votes of the agents in their group, compensating for the agents' costs of voting. Since leaders influence a large number of voters, their decisions have a non-negligible impact on the probability of affecting the electoral outcome and consequently, on the individual instrumental benefit from voting. (Schram 1991; Shachar and Nalebuff 1999) test group based models and provide some empirical support for the mobilization thesis. (see also (Morton 1987, 1991; Uhlaner 1989)). The problem for models of mobilization is that it is not clear how leaders affect the individual behavior of voters.

Models of *group based welfare* consider groups of like-minded individuals whose actions are intended not to maximize their individual utilities, but rather that of the group (see (Kinder and Kiewiet 1979; Markus 1988)). In this case, there is no leader who prescribes behavior as in mobilization models, but instead there is an implicit understanding among agents in the group on appropriate behavior. This idea is developed by (Feddersen and Sandroni 2006), who appeal to (Harsanyi 1980) *Group Rule-Utilitarian Theory* to endogenize the *non-instrumental component* of the utility from voting in a way that preserves the positive relation between closeness of the election and incentive to vote typical of the classic *pivotal-voter models*. In this model, agents derive utility from "doing their part": in the spirit of (Harsanyi 1980), this is understood to mean following the rule that, when followed by all the agents in a given group, would maximize some measure of the group's utility. The outcome is a set of rules for each group, which are mutually optimal (from the point of view of the group) given that individuals follow the rules within their group. Abstention still occurs because for some agents (those with higher costs of voting) the rule prescribes not to vote, since their contribution to increase the group's utility from the election's outcome does not compensate the increase in the group's total cost of voting. (Coate and Conlin 2004) provide some empirical support to the group rule.

The Common Value Setting with Strategic Agents

Feddersen and Pesendorfer (1996, 1997) consider the Condorcet Jury Theorem in a strategic setting and reach a different conclusion than Condorcet. They model a voting problem in an almost common value setting as a game, that is, as a situation where agents interact strategically. The following simple example provides the basic insights of their model: suppose that there are three voters, 1, 2 and 3, and two candidates, A and B . Suppose that voter 1 is an *A-partisan*, meaning he prefers candidate A in all states of the world. Agents 2 and 3 instead prefer candidate A in state S_A and candidate B in state S_B . Let p be the probability of state S_A and $1 - p$ the probability of S_B . Now, suppose that agent 2 is informed, i.e. he knows the state of the world before voting, while agent 3 is not. Finally, suppose that there is no cost of voting and that the election is decided by simple majority rule. In this situation, agent 1 votes for A , agent 3 for B , and agent 2 for A if he observes S_A and for B if he observes S_B . To understand why the uninformed agent 3 votes for B , notice that by doing so the outcome is that A is always selected in state S_A , while B is always selected in state S_B . This is clearly the best outcome for agent 3, and in all states this is also the best outcome for the majority of the population. In fact, if the true state is S_A , A is selected, which is preferred to B by all voters; if the true state is S_B , B is selected, which is preferred to B by two voters out of three.

The uninformed agent in the example votes for B no matter what the prior probability p is, even if p is close to one, that is, even if he is almost sure that A is the right candidate. By voting for B individual 3 counterbalances the *A-partisan's* vote, thereby allowing the informed voter (2) to induce the “right” outcome with probability one. In Condorcet’s argument, if voters vote according to their belief about the state of the world, information is aggregated in a way that induces the right social choice to be made. In this setting, where preferences are only partially aligned, information is aggregated so as to always generate the decision that is preferred by the majority of the population only if uninformed voters vote strategically. In this

setting, if voters vote sincerely, as is assumed by Condorcet, the result that the aggregation of information taking place during an election delivers the “right” social choice does not hold. In Feddersen and Pesendorfer’s framework it is strategic voting that induces the “correct” social choice. This is the major difference with Condorcet’s result, which was driven by the assumption of sincere voting.

Recognizing the strategic incentives that voters may have in an election, Feddersen and Pesendorfer (1998) apply a similar analysis to the unanimity rule in juries. They find that in the context of a common value setting, unanimity voting might result in convicting the innocent more often than other rules because strategic jurors consider the probability of being pivotal (like voter 3 in the example) and make their decisions conditional on being pivotal.

Now suppose there is another voter, 4, who is uninformed and shares the same preferences as 2 and 3. In this setup, even with a zero cost of voting agent 4 would abstain, so that the informed voter is pivotal with probability one. Voter 4’s behavior is known as *strategic abstention*: the act of abstaining by uninformed voters not because voting is costly, but because by doing so they allow the informed voters to be pivotal. By abstaining, uninformed voters in effect delegate the decision to the informed voters. In the example, voter 4’s strategic abstention allows for information equivalence to arise: making voter 4 informed would not change the outcome of the election, as long as 4 strategically abstains so that the informed voter determines the outcome. Notice that this result is in the same spirit as Condorcet’s Jury Theorem’s result: there, having each voter’s belief accurate with a probability slightly higher than 50% or substantially higher than 50% does not make a difference. As long as voters vote according to their belief, the outcome is the same no matter what the underlying belief accuracy is (as long as it is greater than a half). Although for different reasons, in both Condorcet’s setting and Feddersen and Pesendorfer’s model (under some circumstances) the aggregation of information that takes place during an election ensures that the outcome of the election does not vary if the electorate is made more informed.

Future Directions

There are a number of open questions in the literature on voting. As mentioned in earlier sections, at the time of writing there is no universally accepted theory of turnout. In particular, no theory delivers all of the relevant comparative statics observed empirically, that is, the variability of turnout across elections of difference size, importance, and closeness. Broadly speaking, one may say that the behavior of voters is still only partially understood. Whether voters vote based on strategic considerations or vote without regard to how other citizens might be voting is an unsettled issue. Furthermore, individuals' choices in hypothetical and real situations might differ. The act of voting in large elections is almost a hypothetical choice, in that the likelihood that a vote determines the outcome is negligible. An open question then is whether voters choose candidates as they would in a real situation or as they would in a hypothetical situation. As far as empirical work is concerned, little seems to be known of the empirical impact of political ignorance on the outcome of elections. Also, an issue that is both empirical and theoretical and has not been satisfactorily addressed is what types of election rules are best suited for different decisions. Finally, although much has been written on the origins and evolution of democracies, there is no general consensus as to what are the possible reasons of democracies historical evolution. As is clear from this section, many interesting questions in politics have not yet been satisfactorily answered, leaving space for future research.

Bibliography

Primary Literature

- Arrow KJ (1950) A difficulty in the concept of social welfare. *J Polit Econ* 58(4):328–346
- Bentham J (1789) *An introduction to the principles of morals and legislation*. Payne, London
- Besley T, Coate S (1997) An economic model of representative democracy. *Quart J Econ* 112(1):85–114
- Black D (1948) The decisions of a committee using a special majority. *Econometrica* 16(3):245–261
- Blais A (2000) *To vote or not to vote: the merits and limits of rational choice theory*. University of Pittsburgh Press, Pittsburgh
- Borda J (1781) Mathematical derivation of an election system. *Isis* 44(1-2):42–51
- Campbell A, Converse PE, Miller WE, Stokes DE (1960) *The American voter*. Wiley, New York
- Coate S, Conlin M (2004) A group-rule utilitarian approach to voter turnout: theory and evidence. *Am Econ Rev* 95(5):1476–1504
- Condorcet M (1785) *Essai sur l'application de l'analyse de la probabilité des décisions rendues à la pluralité des voix*. L'Imprimerie Royale, Paris. English edition: Baker (1976) Chelsea, New York
- Downs A (1957) *An economic theory of democracy*. Harper, New York
- Duverger M (1972) Factors in a two-party and multiparty system. In: *Party politics and pressure groups*. Crowell Company, New York, pp 23–32
- Dye TR, Zeigler TH (1970) *The irony of democracy, an uncommon introduction to american politics*. Duxbury, Belmont
- Feddersen T, Pesendorfer W (1996) The swing voter's curse. *Am Econ Rev* 86(3):408–424
- Feddersen T, Pesendorfer W (1997) Voting behavior and information aggregation in elections with private information. *Econometrica* 65(3):408–424
- Feddersen T, Pesendorfer W (1998) Convicting the innocent: the inferiority of unanimous jury verdicts. *Am Polit Sci Rev* 92:23–35
- Feddersen T, Sandroni A (2006) Theory A of participation in elections. *Am Econ Rev* 96(4):1271–1282
- Gibbard A (1973) Manipulation of voting schemes. *Econometrica* 41:587–601
- Harsanyi JC (1953) Cardinal utility in welfare economics and the theory of risk-taking. *J Polit Econ* 61:434–435
- Harsanyi JC (1977) *Rational behavior and bargaining equilibrium in games and social situations*. Cambridge University Press, New York
- Harsanyi JC (1980) Rule utilitarianism, rights, obligations and the theory of rational behavior. *Theor Decis* 12(1):115–133
- Hotelling H (1929) Stability and competition. *Econ J* 39(1):41–57
- Kinder DR, Kiewiet DR (1979) Economic discontent and political behavior: the role of personal grievances and collective economic judgements in congressional voting. *Am J Polit Sci* 23(3):495–527
- Ledyard J (1981) The paradox of voting and candidate competition: a general equilibrium analysis. In: Hoorwich G, Quick JP (eds) *Essays in contemporary fields of economics*. Purdue University Press, Lafayette, pp 54–80
- Ledyard J (1984) The pure theory of two candidate elections. *Public Choice* 44(1):7–41
- Markus G (1988) The impact of personal and national economic conditions on the presidential vote: a pooled cross-sectional analysis. *Am J Polit Sci* 32(1):137–154
- Mas-Colell A, Whinston M, Green J (1995) *Microeconomic theory*. Oxford University Press, Oxford
- Morton R (1987) Group majority a model of voting. *Soc Choice Welf* 4(2):117–131

- Morton R (1991) Groups in rational turnout models. *Am J Polit Sci* 35:758–776
- Osborne MJ, Slivinski A (1996) Model a of political competition with citizen-candidates. *Quart J Econ* 111(1):65–96
- Palfrey T, Rosenthal H (1983) A strategic calculus of voting. *Public Choice* 41(1):7–53
- Palfrey T, Rosenthal H (1985) Voter participation and strategic uncertainty. *Am Polit Sci Rev* 79(1):62–78
- Persson T, Tabellini G (2001) Political economics. Press MIT, Boston
- Rawls J (1971) A theory of justice. Harvard University Press, Cambridge
- Riker W, Ordeshook P (1968) A theory of the calculus of voting. *Am Polit Sci Rev* 62:25–42
- Robbins L (1938) Interpersonal comparisons of utility: A comment. *Econ J* 48(192):635–641
- Satterthwaite M (1975) Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. *J Econ Theory* 10:187–217
- Schram A (1991) Voter behavior in economic perspective. Springer, Heidelberg
- Shachar R, Nalebuff B (1999) Follow the leader: theory and evidence on political participation. *Am Econ Rev* 89(3):525–547
- Tullock G (1967) Towards a mathematics of politics. The University of Michigan Press, Ann Arbor
- Uhlener C (1989) Rational turnout: the neglected role of groups. *Am J Polit Sci* 33(2):390–422
- Wolfinger R (1980) Who votes? Yale University Press, New Haven
- Cox GW (1997) Making votes count: strategic coordination in the world's electoral systems. Cambridge University Press, Cambridge
- Dummett M (1984) Voting procedures. Clarendon Press, Oxford
- Duverger M (1959) Political parties: their organization and activity in the modern state. Methuen and Co Ltd, London
- Dworkin R (1981) What is equality? Part 2: equality of resources. *Philos Public Aff* 10:283–345
- Feddersen T (2004) Rational choice theory and the paradox of not voting. *J Econ Perspect* 18(1):99–112
- Katz RS (1997) Democracy and elections. Oxford University Press, Oxford
- McLean I (1987) Public choice: an introduction. Basil Blackwell Inc., New York
- Milnor AJ (1969) Elections and political stability. Little Brown, Boston
- Myerson R (2000) Large poisson games. *J Econ Theory* 94(1):7–45
- Mueller DC (1989) Public choice II: a revised edition of public choice. Cambridge University Press, Cambridge
- Niemi RG, Weisberg HF (eds) (1972) Probability models of collective decision making. Charles Merrill Publishing E company, Columbus
- Ordeshook PC (1986) Game theory and political theory: An introduction. Cambridge University Press, Cambridge
- Ordeshook PC (ed) (1989) Models of strategic choice in politics. The University of Michigan Press, Ann Arbor
- Rawls J (1958) Justice as fairness. *Philos Rev* 67(2):164–194
- Riker WH (ed) (1993) Agenda formation. The University of Michigan Press, Ann Arbor
- Sen AK (1973) On economic inequality. Oxford University Press, Oxford
- Sen AK (2002) Rationality and freedom. The Belknap Press of Harvard University Press, Cambridge
- Tidean N (2006) Collective decisions and voting: the potential for public choice. Ashgate, Burlington
- Tullock G (1998) On voting. Edward Elgar Publishing, Northampton

Books and Reviews

- Arrow KJ (1951) Social choice and individual values. Wiley, New York
- Arrow KJ, Sen AK, Suzumura K (eds) (2002) Handbook of social choice and welfare, vol 1. Elsevier, Amsterdam
- Austen-Smith D, Banks JS (1999) Positive political theory I: collective preference. The University of Michigan Press, Ann Arbor



Voting Procedures, Complexity of

Olivier Hudry
École Nationale Supérieure des
Télécommunications, Paris, France

Article Outline

Glossary
Definition of the Subject
Introduction
Common Voting Procedures
Complexity Results
Further Directions
Bibliography

Glossary

Condorcet winner A candidate is a Condorcet winner if he or she defeats any other candidate in a one-to-one matchup. Such a candidate may not exist; at most, there is only one. Though it could seem reasonable to adopt a Condorcet winner (if any) as the winner of an election, many common voting procedures bypass the Condorcet winner in favor of a winner chosen by other criteria.

Majority relation, strict majority relation In a pairwise comparison method, each candidate is compared to all others, one at a time. If a candidate x is preferred to a candidate y by at least $m/2$ voters (a majority), where m denotes the number of voters, x is said to be preferred to y according to the majority relation. The strict majority relation is defined in a similar way, but with $(m+1)/2$ instead of $m/2$. If there is no tie, the strict majority relation is a tournament, i.e., a complete asymmetric binary relation, called the majority tournament.

Preference, preference aggregation A voter's preference is some relational structure defined

over the set of candidates. Such a structure depends on the chosen voting procedure and usually ranges between a binary relation on one extreme and a linear order on the other. Given a collection, called a profile, of individual preferences defined on a set of candidates, the aggregation problem consists in computing a collective preference summarizing the profile as well as possible (for a given criterion).

Profile A profile $\Pi = (R_1, R_2, \dots, R_m)$ is an ordered collection (or a multiset) of m relations R_i ($1 \leq i \leq m$) for a given integer m . As the relations R_i can be the same, another representation of a profile Π consists in specifying only the q relations R_i which are different, for an appropriate integer q , and the number m_i of occurrences of each relation R_i ($1 \leq i \leq q$): $\Pi = (R_1, m_1; R_2, m_2; \dots; R_q, m_q)$.

Social choice function, social choice correspondence A social choice function maps a collection of individual preferences specified on a set of candidates onto a unique candidate, while a social choice correspondence maps it onto a nonempty set of candidates. This provides a way to formalize what constitutes the most preferred choice for a group of agents.

Voting procedure, voting theory A voting procedure is a rule defining how to elect a winner (single-winner election) or several winners (multiple-winner election) or to rank the candidates from the individual preferences of the voters. Voting theory studies the (axiomatic, algorithmic, combinatorial, and so on) properties of the voting procedures designed in order to reach collective decisions.

Definition of the Subject

One main concern of voting theory is to determine a procedure (also called, according to the context or the authors, *rule*, *method*, *social choice function*, *social choice correspondence*, *system*,

scheme, count, rank aggregation, principle, solution, and so on), for choosing a winner from among a set of candidates, based on the preferences of the voters. Each voter's preference may be expressed as the choice of a single individual candidate or, more ambitiously, a ranked list including all or some of the candidates. Such a situation occurs, obviously, in the field of social choice and welfare (for a broader presentation of the field of social choice and welfare, see, for instance, Aizerman and Aleskerov 1995; Arrow 1963; Arrow and Raynaud 1986; Arrow et al. 2002; Barnett et al. 1995; Barthélemy and Monjardet 1981; Elster and Hylland 1986; Fishburn 1973b; Johnson 1998; Kelly 1987; Moulin 1983; Pattanaik and Salles 1983; and Rowley 1993) and especially of elections (for more about voting theory, see Brams and Fishburn 2002; Dummett 1984; Fischer et al. 2013; Hudry and Monjardet 2010; Hudry et al. 2009; Laslier 2004; Levenglick 1975; Levin and Nalebuff 1995; Merrill and Grofman 1999; Nurmi 1987; Saari 2001; Straffin 1980; Taylor 1995, 2005), but also in many other fields: games, sports, artificial intelligence, spam detection, Web search engines, Internet applications, statistics, and so on.

For a long time, much attention has been paid to the axiomatic properties fulfilled by the different procedures that have been proposed. These properties are important in choosing a procedure, since there is no "ideal" procedure (see next section). More recently, in the late 1980s and the early 1990s, the question has arisen regarding the relative difficulty of computing winners according to a given procedure (for an introduction to computational social choice, see, for instance, Chevaleyre et al. (2007)). The first to study the question may have been J. Orlin in 1981, with a result which remained unpublished (Orlin J 1981, unpublished). The first published results are maybe Y. Wakabayashi's ones in her PhD thesis (Wakabayashi 1986) (see also Wakabayashi 1998), where she deals with the aggregation of binary relations into median orders. The first results on the complexity of the aggregation of orders into median orders seem to be those of (Bartholdi et al. 1989a, b; Hudry 1989). From a practical point of view, it is crucial

to be able to announce the winner in a "reasonable" time. This raises the question of the complexity of the voting procedures, which should be taken into account to the same extent as their axiomatic characteristics.

Below, we will detail the complexity results about several procedures: plurality rule (one-round procedure), plurality rule with runoff (two-round procedure), preferential voting procedure (STV), Borda's procedure, Nanson's procedure, Baldwin's procedure, Condorcet's procedure, Condorcet-Kemeny problem, Slater problem, prudent orders (G. Köhler, K.J. Arrow, and H. Raynaud), maximin procedure (P. B. Simpson), minimax procedure (K. J. Arrow and H. Raynaud), ranked pairs procedure (T. N. Tideman), Copeland's procedure, the top cycle solution (J. H. Smith), the uncovered set solution (P. C. Fishburn, N. Miller), the minimal covering set solution (B. Dutta), Banks's solution, the tournament equilibrium set solution (T. Schwartz), Dodgson's procedure, Young's procedure, approval voting procedure, majority-choice approval procedure (F. Simmons), and Bucklin's procedure.

Section "[Introduction](#)" is devoted to a historic overview and to basic definitions and notation. The common voting procedures are depicted in section "[Common Voting Procedures](#)." Section "[Complexity Results](#)" specifies the complexity of these procedures. Other considerations linked to complexity in the field of voting theory can be found in section "[Further Directions](#)."

Introduction

The Search for a "Good" Voting Procedure, from Borda to Arrow

It is customarily agreed that the search for a "good" voting procedure goes back at least to the end of the eighteenth century, to the works of the chevalier Jean-Charles de Borda (1733–1799) (Borda 1784) and of Marie Jean Antoine Nicolas de Caritat, marquis de Condorcet (1743–1794) (Caritat and marquis de Condorcet 1785), and maybe before (for references upon the historical context, see Arrow 1963; Barthélemy and Monjardet 1981;

Black 1958; Guilbaud 1952; Hägele and Pukelsheim 2001; McLean 1995; McLean and Hewitt 1994; McLean and Urken 1995, 1997; McLean et al. 1995, 2007, and references below).

In the 1770s to the 1780s, J. C. de Borda (1784), a member of the French Academy of Sciences, showed that the plurality rule used at that time by the academy was not satisfactory. Indeed, with such a voting procedure, the winner can be contested by a majority of voters who would all agree to choose another candidate instead of the elected winner. (We may notice that the plurality rule with runoff used in many countries has the same defect.) Borda then suggested another procedure (see below). But, as pointed out by Condorcet in 1784 (Caritat and marquis de Condorcet 1785), the procedure advocated by Borda has the same defect as the one depicted by Borda himself for the plurality rule: a majority of dissatisfied voters could agree to constitute a majority coalition against the candidate elected by Borda's procedure in favor of another candidate. Condorcet then designed a method based on pairwise comparisons. By nature, this method cannot elect a winner who will give rise to a majority coalition against him or her. But, unfortunately, Condorcet's method does not always succeed in finding a winner. If such a winner does exist, i.e., if there exists a candidate who defeats any other candidate in such a pairwise comparison, then this candidate is said to be a *Condorcet winner*; if he or she exists, a Condorcet winner is unique. Actually, the Academy of Sciences decided to adopt Borda's method (until 1803). By the way, notice that, according to I. McLean et al. (2007), "both Ramon Llull (ca 1232–1316) and Nicolaus of Cusa (also known as Cusanus, 1401–1464) made contributions which have been believed to be centuries more recent. Llull promotes the method of pairwise comparison, and proposes the Copeland rule to select a winner. Cusanus proposes the Borda rule, which should properly renamed the Cusanus rule." Despite these historical discoveries, we shall keep the usual names.

After these seminal works, in spite of some works by the Swiss Simon Lhuillier (1750–1840) (Lhuillier 1794) (see also Monjardet 1976), the

Spanish Joseph Isidoro Morales (1797), the French Pierre Claude François Daunou (1761–1840) (Daunou 1803), and Pierre Simon, marquis de Laplace (1749–1827) (Laplace (marquis de) 1795), who, through a slightly different approach, rediscovered Borda's procedure, it seems that the history of the theory of social choice slowed down and almost disappeared until the 1870s or even the 1950s (see Black 1958). In the 1870s, the English reverend Charles Lutwidge Dodgson (1832–1898), also (or maybe better) known as Lewis Carroll, proposed a voting system in which the winner is the candidate who becomes a Condorcet winner with the fewest appropriate changes in voters' preferences (Dodgson 1873, 1874, 1876). Some years later, another English mathematician Edward J. Nanson (1850–1936) (Nanson 1882) on the one hand and the Australian Joseph M. Baldwin (1878–1945) (Baldwin 1926) in 1926 on the other hand slightly modified Borda's method by iteratively eliminating some candidates until only one remains.

None of these methods is utterly satisfactory when we consider the properties which are usually considered desirable. In 1951 and 1963, Kenneth J. Arrow (1963) shows that there does not exist a "good" voting method, with respect to some "reasonable" axiomatic properties; this is known as the famous "impossibility theorem." More precisely, assuming that the preferences of the voters are complete preorders (see below; notice that the set of complete preorders includes the one of linear orders, quite often considered to model the preferences of the voters) and that the result of the voting procedure should also be a complete preorder, K. J. Arrow considered the following properties:

- Unrestricted domain or universality: the voting procedure must be able to provide a result whatever the preferences of the voters are.
- Independence of irrelevant candidates: the collective preference between candidates x and y must depend only on the individual preferences between x and y ; in other words, the collective preference between x and y must remain the same as long as the individual preferences between x and y do not change.

- Unanimity (or Pareto property): if a candidate x is preferred to another candidate y by all the voters, then x must be preferred to y in the collective preference too.

K. J. Arrow showed that if there are at least three candidates (things are much more comfortable with only two candidates!) and at least two voters, the only procedure which satisfies all these conditions at once is the dictatorship, in which one voter (the dictator) imposes his or her preference. Though this impossibility theorem ruins the hope to design a voting procedure fulfilling the usual desirable properties, several procedures have been suggested since this date; we shall describe some of them below. Among the ways to escape Arrow's impossibility theorem, we find the following:

- The definition of other axiomatic systems which would lead to voting procedures which would not be dictatorship
- The restriction of the individual preferences to more constrained domains
- Adapting the result, when this one is not satisfactory with respect to the required axiomatic properties, into a result fulfilling these properties and fitting the genuine result as well as possible, for some criterion which must be defined

The main questions associated with the first possibility are "given some axiomatic properties, what are the voting procedures satisfying these properties?" or, conversely, "given a voting procedure, what is the proper axiomatic system characterizing this procedure?"; we will not consider this direction here. The second possibility will be illustrated below, by the restriction of the individual preferences to *single-peaked preferences*. The third direction was followed by J. G. Kemeny in 1959 (Kemeny 1959), when he studied the aggregation of complete preorders into a *median complete preorder* (see below). Notice that the median procedure is also attributed to Condorcet; in the sequel, we will refer the search for a median linear order as the Condorcet-Kemeny problem (as other people rediscovered this problem or some of its

variants, Monjardet 1990); the problem is also known under other names – Charon and Hudry (2007). Another related problem, dealing with the *majority tournament* (see below), is the one stated explicitly by P. Slater in 1961 (Slater 1961) of fitting a tournament into a linear order at minimum distance. This is one of the so-called tournament solutions (see Laslier 1997; Moulin 1986); for references on tournaments, see also McKey (2013), Moon (1968), Reid (2004), and Reid and Beineke (1978), of which the aim is to determine a winner from a tournament. Besides the Slater solution, we shall describe other tournament solutions, such as the top cycle, the uncovered set, the minimal covering set, Banks's solution, and the tournament equilibrium set. The other common tournament solutions are polynomial, or their complexities are not completely known (see Hudry 2009 for more details).

Definitions, Notation, and Partially Ordered Sets Used to Model Preferences

Let us assume that we are dealing with m voters who must choose between n candidates denoted $x_1, x_2, \dots, x_n$ or $x, y, z \dots$; X denotes this set of candidates; in the following, we suppose that n is large enough. A binary relation R defined on X is a subset of $X \times X = \{(x, y) : x \in X \text{ and } y \in X\}$. We use the notation xRy instead of $(x, y) \in R$ and $x\bar{R}y$ instead of $(x, y) \notin R$.

It is customary to represent the preferences $R_i (1 \leq i \leq m)$ of m voters as an ordered collection (or a multiset) $\Pi = (R_1, R_2, \dots, R_m)$, called the *profile* of the m binary relations. Another representation of the preferences of the voters exists. Since these preferences may be the same for two different voters, we may consider only the different relations $R_1, R_2, \dots, R_q$ arising from the opinions of the voters, where q denotes this number of different opinions. In this way, we combine all the voters sharing the same opinion. Last, if $m_i (1 \leq i \leq q)$ denotes the number of voters sharing $R_i (1 \leq i \leq q)$ as their preference (notice the equality $\sum_{i=1}^q m_i = m$), we associate this number m_i of occurrences with each type R_i of relations to describe Π . Then

Π can be described as the set of such pairs: $\Pi = \{(R_1, m_1); (R_2, m_2); \dots; (R_q, m_q)\}$. Such a representation does represent the data more compactly when m is large with respect to n . Usually, the complexity results are the same for the two representations, because their proofs stand even if m is bounded by a polynomial in n , and in this case, there is no qualitative difference between the two representations. With respect to the results stated in section “Complexity Results,” there would be a difference only for the L^{NP} -completeness results of Theorems 5, 14, and 15 (which then should be replaced only by NP-hardness results). So, in the sequel, we shall consider only the first representation. Moreover, we shall assume that the preference of each voter i ($1 \leq i \leq m$) is given by a binary relation R_i defined on X and that R_i is described by its characteristic vector (see below), which requires n^2 bits. So the size of the data set is about mn^2 .


We will consider two kinds of elections: we want to elect one candidate, or we want to rank all of them into a partially ordered set (or *poset*). One of the most common posets is the structure of linear order. Other posets can be defined from the following basic properties:

Reflexivity	$\forall x \in X, xRx$
Irreflexivity	$\forall x \in X, x\bar{R}x$
Antisymmetry	$\forall (x, y) \in X^2, (xRy \text{ and } x \neq y) \Rightarrow y\bar{R}x$
Asymmetry	$\forall (x, y) \in X^2, xRy \Rightarrow y\bar{R}x$
Transitivity	$\forall (x, y, z) \in X^3, (xRy \text{ and } yRz) \Rightarrow xRz$
Completeness	$\forall (x, y) \in X^2, \text{ with } x \neq y, xRy, \text{ or (inclusively) } yRx$


By combining the above properties, we may define different types of binary relations (see, for instance, Barthélemy and Monjardet 1981; Caspard et al. 2007; Fishburn 1973a, 1985). As a binary relation R defined on X is the same as the directed graph, $G = (X, R) : (x, y)$ is an arc (i.e., a directed edge; similarly, a directed cycle will be called a circuit in the sequel) of G if and only if we have xRy ; we illustrate these types with graph theoretic examples (for basic references on graph theory, see Bang-Jensen and Gutin 2001 or Berge 1985). It is often possible to define reflexive or irreflexive versions of the following ordered

structures. But as reflexivity and irreflexivity do not matter for complexity results (see Hudry 2008), we give below only one version among these two possibilities.


- A *partial order* is an antisymmetric (if reflexive) or asymmetric (if irreflexive) and transitive binary relation (see Fig. 1); $\mathcal{O}$ will denote the set of the partial orders defined on X .
- A *linear order* is a complete partial order (see Fig. 2); $\mathcal{L}$ will denote the set of the linear orders defined on X . If L denotes a linear order defined on X , we will represent L as $x_{\sigma(1)} > x_{\sigma(2)} > \dots > x_{\sigma(n)}$ for some appropriate permutation σ , with the agreement that the notation $x_{\sigma(i)} > x_{\sigma(i+1)}$ (for $1 \leq i < n$) means that $x_{\sigma(i)}$ is preferred to $x_{\sigma(i+1)}$ according to L , the relationship between the other elements of X being involved by transitivity. The element $x_{\sigma(1)}$ will be called *the winner of L*.
- A *tournament* is a complete and asymmetric binary relation (see Fig. 3); $\mathcal{T}$ will denote the set of the tournaments defined on X ; notice that a transitive tournament is a linear order and conversely. As a tournament may contain circuits, they are usually not considered as an appropriate structure to represent the collective preference that we seek. Tournaments will be




Voting Procedures, Complexity of, Fig. 1 A partial order




Voting Procedures, Complexity of, Fig. 2 A linear order. The partial order of Fig. 1 is not a linear order, for instance, because the vertices a and d are not compared




Voting Procedures, Complexity of, Fig. 3 A tournament



Voting Procedures, Complexity of, Fig. 5 A complete preorder



Voting Procedures, Complexity of, Fig. 4 A preorder



Voting Procedures, Complexity of, Fig. 6 An acyclic relation. Its transitive closure is the partial order of Fig. 1

used below to summarize individual preferences in the so-called *majority tournament*.

- A *preorder* is a reflexive and transitive binary relation (see Fig. 4); $\mathcal{P}$ will denote the set of the preorders defined on X .
- A *complete preorder* is a reflexive, transitive, and complete binary relation (see Fig. 5); $\mathcal{C}$ will denote the set of the complete preorders defined on X .
- An *acyclic relation* (we should rather say *without circuit*, but *acyclic* is the usual term) is a relation R of which the associated graph G is without any circuit (see Fig. 6); $\mathcal{A}$ will denote the set of the acyclic relations defined on X .

As stated above, it is possible to obtain other structures by adding or removing reflexivity or irreflexivity from the above definitions. In fact, the distinction between reflexive and irreflexive relations is not relevant from the complexity point of view: the results remain the same. Thus, in the following, we do not take reflexivity or irreflexivity into account.

Other types of posets exist, obtained by considering the asymmetric parts of the previous relations (for instance, a *weak order* is sometimes defined as the asymmetric part of a complete preorder) or by adding extra structural properties (for instance, to define interval orders or semi-orders from partial orders). We will not consider these posets here; the interested reader is referred to Caspard et al. (2007) for their definitions and to Hudry (2008) for some complexity results about them. On the contrary, we will pay attention to generic binary relations without any particular property; the set of the binary relations will be noted $\mathcal{R}$. We may notice several inclusions between these sets, especially the following ones: $\forall Z \in \{\mathcal{A}, \mathcal{C}, \mathcal{L}, \mathcal{O}, \mathcal{P}, \mathcal{R}, \mathcal{T}\}, \mathcal{L} \subseteq \mathcal{Z} \subseteq \mathcal{R}$; in other words, a linear order can be considered as a special case of any one of the other types, and any type is a special case of binary relation.

To conclude this section, we summarize the notation for these sets in Table 1.

Voting Procedures, Complexity of, Table 1 Meaning of the notation $\mathcal{A}$, $\mathcal{C}$, $\mathcal{L}$, $\mathcal{O}$, $\mathcal{P}$, $\mathcal{R}$, $\mathcal{T}$

$\mathcal{A}$: acyclic relation	$\mathcal{L}$: linear order	$\mathcal{P}$: preorder	$\mathcal{T}$: tournament
$\mathcal{C}$: complete preorder	$\mathcal{O}$: partial order	$\mathcal{R}$: binary relation	

Common Voting Procedures

In this section, we describe the main voting procedures, i.e., plurality rule (one-round procedure), plurality rule with runoff (two-round procedure), preferential voting procedure (STV), Borda's procedure, Nanson's procedure, Baldwin's procedure, Condorcet's procedure, Condorcet-Kemeny problem, Slater problem, prudent orders (G. Köhler, K. J. Arrow, and H. Raynaud), maximin procedure (P. B. Simpson), minimax procedure (K. J. Arrow and H. Raynaud), ranked pairs procedure (T. N. Tideman), Copeland's procedure, the top cycle solution (J. H. Smith), the uncovered set solution (P. C. Fishburn, N. Miller), the minimal covering set solution (B. Dutta), Banks's solution, the tournament equilibrium set solution (T. Schwartz), Dodgson's procedure, Young's procedure, approval voting procedure, majority-choice approval procedure (F. Simmons), and Bucklin's procedure (see also Brams and Fishburn 2002). For other tournament solutions, see Laslier (1997) and Moulin (1986).

Plurality Rule, Plurality Rule with Runoff, and Preferential Voting Procedure

One of the easiest voting procedures by which to elect one candidate as a winner is the *plurality rule* (also called *one-round procedure* or *relative majority*, or sometimes *first-past-the-post*, or *winner-takes-all*, or also *majoritarian voting*...; see Inada (1969)). In this procedure, each voter gives 1 point to his or her favorite candidate (so it is not necessary to know the preferences of the voters on the whole set of candidates). The candidate who gains the maximum number of points is the winner.

This procedure belongs to the family of *scoring procedures*. In such a procedure, a *score vector* $(s_1, s_2, \dots, s_n)$ is fixed independently of the voters, with $s_1 \geq s_2 \geq \dots \geq s_n$. For each voter, a candidate x receives s_i points if x is ranked

at the i th position by the considered voter. The *score* of x is the total number of points that x received. The winner is the candidate with the maximum score. If the aim is to rank the candidates, we may also sort them according to decreasing scores and then consider the linear extensions of the complete preorder provided by this sorting. (Another possibility would be to apply the procedure $n - 1$ times, after having removed the winner of the current iteration.) For the plurality rule, the score vector is $(1, 0, 0, \dots, 0)$.

There are two rounds in the *plurality rule with runoff*, also called *two-round* (or *two-ballot*) *procedure*. The first round is like the plurality rule described above. At the end of this first step, if a candidate has gained at least $(m + 1)/2$ points (the strict majority), he or she is the winner. Otherwise, the two candidates with the maximum numbers of points remain for a second round, and the others are removed from the election. Then the plurality rule is applied again but only to the remaining two candidates. This method is designed to elect only one winner. If we want to obtain k winners, it is sufficient to apply it k times. Notice that the repetition of a given procedure always makes it possible to elect several winners.

A generalization of these procedures consists in performing a given number of rounds. For each round, each voter gives 1 point for his or her favorite candidate. If, at the end of a round, there is a candidate who has gained at least $(m + 1)/2$ points, he or she is the winner. Otherwise, the candidates with the lowest numbers of points are eliminated from the competition; the number of candidates eliminated at each round depends on the number of rounds but is such that only two candidates will remain for the last round. The winner of this last round is the winner of the election. Special cases are the plurality rule with runoff, described above, and the one with at most $n - 1$ rounds. For this latter case, exactly one candidate is removed at each round. This variant,

in which the candidate who is the least often ranked at the first position is removed, is also known as *preferential voting* (or *preference voting*) or as *single transferable vote* (STV) or as *instant-runoff voting* (IRV).

Other variants of these voting procedures may be defined by the successive eliminations of the losers. For instance, the candidate who is most often ranked last is removed, and we iterate this process while there remain at least two candidates.

Borda's Procedure and Some Variants (Nanson's and Baldwin's Procedures)

As related above, Borda considered the plurality rule unsatisfactory because the winner can be contested by a majority of voters who would all agree to choose another candidate instead of the elected winner (the same holds for the plurality rule with runoff, see Example 1). Borda suggested another procedure in which the voters rank the n candidates according to their preferences. For each voter, the candidate who is ranked first is given $n - 1$ points, a candidate ranked second is given $n - 2$ points, and so on: more generally, a candidate ranked at the i th position is given $n - i$ points. Then, all these points are summed up for each candidate: this sum is the *Borda score* s_B of the candidate. The candidate with a maximum Borda score is the *Borda winner*. So Borda's procedure is also a scoring procedure, of which the score vector is $(n - 1, n - 2, \dots, 1, 0)$.

Using Borda's procedure, one can easily obtain a ranking of all the candidates. A first possibility consists, as for the plurality rule or the plurality rule with runoff, in iterating Borda's procedure $n - 1$ times, after having removed the Borda winner of the current iteration. But we may also apply Borda's procedure only once and then rank the candidates according to the decreasing values of their Borda scores. This gives a complete pre-order. Any linear extension of this complete pre-order can be considered as the collective ranking according to Borda's procedure. Notice that these two possibilities do not necessarily provide the same rankings (see Example 1).

Several variants of Borda's procedure have been studied. For instance, we may apply other score vectors: instead of $n - i$ points given to the

candidate ranked at the i th position, we may choose to assign other values. For instance, given an integer k , we may credit k points to the candidate ranked at the first position, $k - 1$ points to the second, and so on, through the candidate ranked at the k th position, who gains 1 point, after which the following candidates gain nothing (the shape of the score vector is $(k, k - 1, \dots, 1, 0, \dots, 0)$). For $k = 1$, this system is the plurality rule. For $k \geq n - 1$, this system gives the same results as Borda's procedure.

Other systems are based on Borda's procedure but with the elimination of some candidates. *Nanson's procedure* (Nanson 1882) modifies Borda's procedure by eliminating the candidates whose Borda scores are below the average Borda score and by repeating the computations of the Borda scores with respect to the remaining candidates after these eliminations, until there remains only one candidate. Another variant of Borda's procedure is one suggested by Baldwin in 1926 (Baldwin 1926); as in Nanson's procedure, candidates are iteratively removed from the election. But, in *Baldwin's procedure*, only one candidate is removed at each iteration, the one whose Borda score is the lowest.

Condorcet's Procedure

Condorcet designed a method based on pairwise comparisons. More precisely, for each candidate x and each candidate y with $x \neq y$, we compute the number m_{xy} that we will call the *pairwise comparison coefficient* below, of voters who prefer x to y . Then x is considered as better than y if a majority of voters prefers x to y , i.e., if we have $m_{xy} > m_{yx}$. This defines the (*strict*) *majority relation* T : $xTy \Leftrightarrow m_{xy} > m_{yx}$. In some cases, there exists a *Condorcet winner*, i.e., a candidate x defeating any other candidate: $\forall y \neq x, m_{xy} > m_{yx}$. If there exists a Condorcet winner, then he or she is unique. It may even happen that T is a linear order and allows us to rank all the candidates. Such is the case for the following example (due to B. Monjardet, private communication), which illustrates the previous voting procedures.

Example 1 Assume that $m = 13$ voters must rank $n = 4$ candidates x, y, z , and t . Suppose the preferences of the voters are given by the following linear orders:

- The preferences of two voters are $x > y > z > t$.
- The preference of one voter is $y > z > x > t$.
- The preference of one voter is $y > z > t > x$.
- The preferences of four voters are $z > y > x > t$.
- The preferences of five voters are $t > x > y > z$.

According to the plurality rule, t is the winner with 5 points (2 points for x and for y , 4 points for z). According to the plurality rule with runoff, z is the winner (the four voters who voted for x or y prefer z to t). The Borda scores s_B of the candidates are:

- $s_B(x) = 2 \times 3 + 5 \times 2 + 5 \times 1 + 1 \times 0 = 21$;
- $s_B(y) = 2 \times 3 + 6 \times 2 + 5 \times 1 + 0 \times 0 = 23$;
- $s_B(z) = 4 \times 3 + 2 \times 2 + 2 \times 1 + 5 \times 0 = 18$;
- $s_B(t) = 5 \times 3 + 0 \times 2 + 1 \times 1 + 7 \times 0 = 16$.

So the winner according to Borda's procedure is y , and the ranking of the four candidates is $y > x > z > t$ (notice that if, while there are at least two candidates, we apply the variant consisting in removing the winner and reapplying Borda's procedure, then we obtain the orders $y > x > z > t$ and $y > z > x > t$ as the possible rankings; the distance to the ranking provided by one application of Borda's procedure may be much more important). In Nanson's procedure, as the Borda scores of z and t are below the average (which is equal to 19.5), z and t are removed and only x and y remain. Then the Borda scores of x and y for the second step become $s_B(x) = 7$ and $s_B(y) = 6$; hence, x is the winner according to Nanson's procedure. In Baldwin's procedure, t is first removed. The Borda scores of x, y , and z become

$s_B(x) = 14$, $s_B(y) = 15$, and $s_B(z) = 10$. Thus, z is removed and the remaining computations are as in Nanson's procedure: here also x is the winner (but it may happen that Nanson's procedure and Baldwin's procedure do not provide the same winners).

Let us now compute the pairwise comparison coefficients m_{xy} necessary to apply Condorcet's procedure:

- $m_{xy} = 2 + 5 = 7$; $m_{yx} = 1 + 1 + 4 = 6$;
- $m_{xz} = 2 + 5 = 7$; $m_{zx} = 1 + 1 + 4 = 6$;
- $m_{xt} = 2 + 1 + 4 = 7$; $m_{tx} = 1 + 5 = 6$;
- $m_{yz} = 2 + 1 + 1 + 5 = 9$; $m_{zy} = 4$;
- $m_{yt} = 2 + 1 + 1 + 4 = 8$; $m_{ty} = 5$;
- $m_{zt} = 2 + 1 + 1 + 4 = 8$; $m_{tz} = 5$.

The bold values are the ones greater than the majority $(m + 1)/2$. These values show that the majority relation, here, is a linear order, namely, $x > y > z > t$.

By the way, this example shows the importance of the voting procedure, as already noticed by Borda (see Mascart 1919): four procedures and four different winners; so all four of our candidates may claim to be the winners of the election. . . . It is often the case that different procedures provide different winners.

Median Orders, Condorcet-Kemeny Problem, and Slater Problem

As Condorcet himself discovered, his procedure may lead to a majority relation which is not transitive. The simplest example in this respect is one with $n = 3$ candidates x, y , and z and with $m = 3$ voters whose preferences are, respectively, $x > y > z$, $y > z > x$, and $z > x > y$. It is easy to verify that the majority relation T is defined by xTy , yTz , and zTx , hence a lack of transitivity. If there is no tie (which is necessarily the case if m is odd, since the individual preferences of the voters are assumed to be linear orders), then T is a tournament, called the *majority tournament*. We may or may not weigh T if we want to take the intensity of the preferences into account. If T is not weighted, the search for a winner or for a ranking of the candidates leads to the definition of *tournament solutions* (see Laslier (1997) for a comprehensive

study of these and Hudry (2009) for a survey of their complexities, some of which are given below). In the Condorcet-Kemeny problem, T is weighted, and the aim is to compute a linear order or, more generally, a poset fitting T “as well as possible.”

To specify what “as well as possible” means, we use the *symmetric difference distance* δ defined, for two binary relations R and S defined on X , by

$$\delta(R, S) = |\{(x, y) \in X^2 : (xRy \text{ and } x\bar{S}y) \text{ or } (x\bar{R}y \text{ and } xSy)\}|.$$

This quantity $\delta(R, S)$ measures the number of disagreements between R and S . Though it is possible to consider other distances, δ is widely used and is appropriate for many applications. J. P. Barthélemy (1979) shows that δ satisfies a number of naturally desirable properties. J. P. Barthélemy and B. Monjardet (1981) recall that $\delta(R, S)$ is the Hamming distance between the characteristic vectors (see below) of R and S and point out the links between δ - and the L_1 -metric or the square of the Euclidean distance between these vectors (see also Monjardet 1979, 1990).

So for a profile $\Pi = (R_1, R_2, \dots, R_m)$ of m relations, we can define the *remoteness* $\Delta(\Pi, R)$ between a relation R and the profile Π by

$$\Delta(\Pi, R) = \sum_{i=1}^m \delta(R, R_i).$$

The remoteness $\Delta(\Pi, R)$ measures the total number of disagreements between Π and R (see Hudry 2013a) for the study of the complexity for other kinds of remoteness).

Our aggregation problem can now be seen as a combinatorial optimization problem: given a profile Π , determine a binary relation R^* minimizing Δ over one of the sets $\mathcal{A}, \mathcal{C}, \mathcal{L}, \mathcal{O}, \mathcal{P}, \mathcal{R}, \mathcal{T}$. Such a relation R^* will be called a *median relation* of Π (Barthélemy and Monjardet 1981). According to the number m of relations of the profile, and according to the properties assumed for the relations belonging to Π or required from the median relation, we get many combinatorial problems. We note them as follows:

Problems $P_m(\mathcal{Y}, \mathcal{Z})$. For $\mathcal{Y}$ belonging to $\{\mathcal{A}, \mathcal{C}, \mathcal{L}, \mathcal{O}, \mathcal{P}, \mathcal{R}, \mathcal{T}\}$ and $\mathcal{Z}$ also belonging to $\{\mathcal{A}, \mathcal{C}, \mathcal{L}, \mathcal{O}, \mathcal{P}, \mathcal{R}, \mathcal{T}\}$, for a positive integer m , $P_m(\mathcal{Y}, \mathcal{Z})$ denotes the following problem: given a finite set X of n elements and given a profile Π of m binary relations all belonging to $\mathcal{Y}$, find a relation R^* belonging to $\mathcal{Z}$ minimizing Δ over $\mathcal{Z}$: $\Delta(\Pi, R^*) = \min\{\Delta(\Pi, R) \text{ for } R \in \mathcal{Z}\}$.

With this notation, the initial problem, possibly considered by Condorcet, is $P_m(\mathcal{L}, \mathcal{L})$, consisting in aggregating m linear orders into a median linear order. Similarly, the problem considered by J. G. Kemeny (1959), consisting in aggregating m complete preorders into a median complete preorder, is $P_m(\mathcal{C}, \mathcal{C})$. A *Condorcet-Kemeny winner* is the winner of any median linear order (for references about the Condorcet-Kemeny problem, see, e.g., Charon and Hudry 2007; Jünger 1985; Monjardet 2008a, b; Reinelt 1985; and references therein).

We may easily state the problems $P_m(\mathcal{Y}, \mathcal{Z})$ as 0–1 linear programming problems (see Barthélemy and Monjardet 1981; Charon and Hudry 2007; Hudry 1989, 2008; Wakabayashi 1986 for instance) for any profile $\Pi = (R_1, R_2, \dots, R_m)$ of m binary relations R_i ($1 \leq i \leq m$) all belonging to $\mathcal{Y}$. To this end, consider the characteristic vectors $r^i = (r_{xy}^i)_{(x,y) \in X^2}$ of the relations R_i ($1 \leq i \leq m$) defined by $r_{xy}^i = 1$ if $xR_i y$ and $r_{xy}^i = 0$ otherwise and similarly the characteristic vector $r = (r_{xy})_{(x,y) \in X^2}$ of any binary relation R . Then, after some computations, we obtain $\Delta(\Pi, R) = C - \sum_{(x,y) \in X^2} \alpha_{xy} r_{xy}$, where $C = \sum_{i=1}^m \sum_{(x,y) \in X^2} r_{xy}^i$ is a constant and with $\alpha_{xy} = 2 \sum_{j=1}^m r_{xy}^j - m$. Notice that, for

the problem $P_m(\mathcal{L}, \mathcal{L})$, the sum $\sum_{i=1}^m r_{xy}^i$ is equal for

$x \neq y$ to the pairwise comparison coefficient m_{xy} introduced above. So, because of the equality

$$m = \sum_{i=1}^m r_{xy}^i + \sum_{i=1}^m r_{yx}^i, \text{ we get } \alpha_{xy} = m_{xy} - m_{yx}:$$

α_{xy} measures the number of voters who prefer x to y minus the number of voters who prefer y to x . More generally, α_{xy} is equal to twice the gap between the number of voters who prefer x to

y and the majority. It is a non-positive or non-negative integer with the same parity as m . To obtain a 0–1 linear programming statement of $P_m(\mathcal{Y}, \mathcal{Z})$, it is then sufficient to express the constraints defining the set $\mathcal{Z}$, which is easy for the posets described above. For instance, the transitivity of R can be expressed by the following inequalities:

$$\forall(x, y, z) \in X^3, 0 \leq r_{xy} + r_{yz} - r_{xz} \leq 1.$$

As stated above, it is also common to represent a preference R defined on X by a graph. The properties of the graph are the properties of R : it can be antisymmetric, complete, transitive, and so on. Similarly, the profile $\Pi = (R_1, R_2, \dots, R_m)$ can also be represented by a directed, weighted, complete, and symmetric graph $G_\Pi = (X, U_X)$: its set of vertices is X , and G_Π contains all the possible arcs except the loops, i.e., $U_X = X \times X - \{(x, x) \text{ for } x \in X\}$. (The loops would be associated with reflexivity; this property, as well as irreflexivity, has no impact on the complexity status of the studied problems, hence this simplification.) The weights of the arcs (x, y) give the intensity of the preference for x over y . The computations above lead us to assign α_{xy} as the weight of any arc (x, y) of G_Π . With this choice, minimizing $\Delta(\Pi, R)$ is the same, from the graph theoretic point of view, as drawing from G_Π a subset of arcs with a maximum total weight and satisfying the structural properties required from R . Notice that characterizations of the graphs that we can associate with profiles Π have been provided by different authors (see Charon and Hudry 2007; Debord 1987a, b; Erdős and Moser 1964; Hudry 2008; Mc Garvey 1953; Stearns 1959); the construction of the profiles can be done in polynomial time, which allows us to study the complexities of the problems $P_m(\mathcal{Y}, \mathcal{Z})$ through their graph theoretic representations G_Π .

Example 2 illustrates these considerations.

Example 2 Assume that $m = 9$ voters must rank $n = 4$ candidates x, y, z , and t . Assume, also, that the preferences of the voters are given by the following linear orders:

- The preferences of three voters are $x > y > z > t$.
- The preferences of two voters are $y > t > z > x$.
- The preference of one voter is $t > z > x > y$.
- The preference of one voter is $x > z > y > t$.
- The preference of one voter is $t > y > x > z$.
- The preference of one voter is $z > t > y > x$.


The quantities m_{xy} involved in Condorcet's procedure are the following, where the bold values, again, denote those greater than the majority $(m + 1)/2$:

- $m_{xy} = \mathbf{5}$; $m_{yx} = 4$;
- $m_{xz} = \mathbf{5}$; $m_{zx} = 4$;
- $m_{xt} = 4$; $m_{tx} = \mathbf{5}$;
- $m_{yz} = \mathbf{6}$; $m_{zy} = 3$;
- $m_{yt} = \mathbf{6}$; $m_{ty} = 3$;
- $m_{zt} = \mathbf{5}$; $m_{tz} = 4$.

Here, the majority relation is not a linear order but the tournament of Fig. 7. Figure 7 also displays the graph G_Π summarizing the data. We may observe that the majority tournament is given by the arcs of G_Π with a positive weight.

For these data, it is not too difficult to verify that there is only one median linear order, which is $x > y > z > t$, this order keeps all the positive weights except that of the arc (t, x) , of which the weight is the minimum, while the arcs with positive weights do not define a linear order. Hence, x is the only Condorcet-Kemeny winner.

Attention has also been paid to $P_1(\mathcal{T}, \mathcal{L})$, i.e., the approximation of a tournament (which can be the majority tournament of the election) by a linear order at minimum distance. This problem is also known as *Slater problem*, since P. Slater explicitly stated it in 1961 (Slater 1961); see also Charon and Hudry (2007) for a survey of this problem. A linear order L^* at minimum distance (with respect to the symmetric difference distance) from the tournament T constituting the considered instance of $P_1(\mathcal{T}, \mathcal{L})$ is called a *Slater order* of T : $\delta(T, L^*) = \min_L \delta(T, L)$. This minimum distance is usually called the *Slater index* $i(T)$ of T : $\delta(T, L^*) = i(T)$. A *Slater winner* of T is the winner of any Slater order of T .



Voting Procedures, Complexity of, Fig. 7 The majority tournament (left) and the graph G_Π (right) associated with the data of Example 2 and weighted by the quantities $\alpha_{xy} = m_{xy} - m_{yx}$

Prudent Orders, Maximin Procedure, Minimax Procedure, and Ranked Pairs Procedure

Other procedures are based on the quantities m_{xy} defined above.

Such is the case for *prudent orders* proposed by G. Köhler (1978) and by K. J. Arrow and H. Raynaud (1986). Given the pairwise comparison coefficients m_{xy} , let us define the *cut relation* $R_{>t}$ for any integer t with $-m \leq t \leq m$ by the following:

$$\text{for } x \in X \text{ and } y \in X \text{ with } x \neq y, xR_{>t}y \\ \Leftrightarrow m_{xy} - m_{yx} > t.$$

Then define a threshold $t_{\min}$ by

$$t_{\min} = \min \{t \text{ with } -m \leq t \leq m \text{ and such that } R_{>t} \text{ contains no circuit}\}.$$

A *prudent order* is any linear order which contains $R_{>t_{\min}}$. A winner according to the procedure of the prudent orders is the winner of any prudent order of the profile being considered. For Example 2, we have $t_{\min} = 1$; with respect to the graph G_Π of Fig. 7, $R_{>t_{\min}}$ contains only the arcs (y, z) and (y, t) . There are eight prudent orders, including $x > y > z > t$ or $y > t > z > x$, but none of them has z or t as its winner: x and y are the winners according to the prudent order procedure.

The *maximin procedure*, proposed by P. B. Simpson (1969), is also based on the quantities m_{xy} . In this procedure, for each candidate x , we consider the worst performance $W(x)$ in the pairwise comparisons: $W(x) = \min_{y \neq x} m_{xy}$. The winner is any candidate x^* with a maximum worst performance: $W(x^*) = \max_{x \in X} W(x)$. We may also use these quantities to sort the candidates in decreasing order with respect to the quantities W to obtain a collective preference over the whole set of candidates (as usual, another possibility is to remove the winner and to iterate the process). For instance, applied to Example 2, we obtain $W(x) = 4$, $W(y) = 4$, $W(z) = 3$, and $W(t) = 3$: x and y are the Simpson winners.

The *minimax procedure*, in which the previous roles of “min” and “max” are switched and in which the candidates are sorted in increasing order with respect to the obtained maxima, has been proposed by K. J. Arrow and H. Raynaud (1986). For Example 2, the winners would be x , z , and t .

In 1987, T. N. Tideman introduced the *ranked pairs procedure* (Tideman 1987). Here, we first sort the quantities m_{xy} decreasingly. Then we scan them in this order to build a linear order L for the collective preference. If m_{xy} is the current quantity that we scan, then we definitely set xLy if such a decision is compatible with the previous ones. We do so until a linear order is completely defined. For instance, for Example 2, we first fix yLz and yLt since m_{yz} and m_{yt} are maxima and z and t

cannot be winners. Then we must choose between adding xLy or xLz or tLx or zLt , since m_{xy} , m_{xz} , m_{tx} , and m_{zt} are equal. We cannot add all of them simultaneously, because of incompatibilities with yLz or yLt already fixed. If we choose to add xLy , xLz , and zLt , we obtain the linear order $x > y > z > t$, and x is a winner. If we choose to add tLx and zLt , we obtain the linear order $y > z > t > x$, and y is a winner. Here again, x and y are the winners.

Tournament Solutions

The procedures described in this section apply to tournaments, which can function as the majority tournament of an election. They deal with unweighted tournaments, but some of them can be extended to weighted tournaments, as in the case of the solution designed by P. Slater (see above).

Number of Wins: Copeland's Procedure

The procedure designed by A. H. Copeland in 1951 (Copeland 1951) is also based on the m_{xy} 's. For any two different candidates x and y , we set $C(x, y) = 1$ if $m_{xy} > m_{yx}$, $C(x, y) = 0.5$ if $m_{xy} = m_{yx}$, and $C(x, y) = 0$ if $m_{xy} < m_{yx}$. The Copeland score $C(x)$ of x is $C(x) = \sum_{y \neq x} C(x, y)$. The Copeland winner is any candidate with a maximum Copeland score (here also, we may rank the candidates according to decreasing Copeland scores).

The application of Copeland's procedure to Example 2 gives $C(x) = 2$, $C(y) = 2$, $C(z) = 1$, and $C(t) = 1$. Here, x and y are the Copeland winners.

The Copeland score has a meaning from a graph theoretic point of view. If we assume that there is no tie, the majority relation is a tournament (the majority tournament). Then the Copeland score of a candidate x is also the out-degree of the vertex associated with x . Copeland's procedure is one of the common tournament solutions (see Laslier 1997). It can be extended to weighted tournaments by sorting the vertices x according to the sum $\sum_{y \neq x} m_{xy}$ of the pairwise comparison coefficients m_{xy} ; this variant leads to the same ranking as Borda's procedure.

Top Cycle: Smith's Solution

Another tournament solution is the so-called top cycle (also called the Smith set) (Smith 1973). Any directed graph G can be decomposed into its strongly connected components. We may then define another graph H derived from G . The vertices of H are associated with the strongly connected components of G . Let x (respectively y) be a vertex of G associated with the strongly connected components C_x (respectively C_y) of G . There will exist an arc (x, y) from x to y if there is at least one arc in G from C_x to C_y (notice that, in this case, all the arcs between C_x and C_y are from C_x toward C_y). Then H does not contain any circuits. Moreover, if G is a tournament, H is a linear order and admits a winner (but H contains only one vertex if G is strongly connected). The top cycle of G is the strongly connected component of G associated with the winner of H .

The top cycle solution, when applied to a tournament T , consists of considering all vertices of the top cycle of T as the winners of T . For instance, for the tournament of Fig. 7, which is strongly connected, the top cycle contains the four vertices x , y , z , and t , which are the four winners according to this solution.

Uncovered Set: Fishburn's and Miller's Solution

A refinement of the top cycle is provided by the set of uncovered vertices. Given a tournament T and two vertices x and y of T , we say that x covers y if (x, y) is an arc of T and if, for any arc (y, z) , (x, z) is also an arc of T . (In other words, considering T as a majority tournament, x beats y and any vertex beaten by y is also beaten by x .) A vertex is said to be uncovered if no other vertex covers it. The uncovered set of T is noted $UC(T)$. Adopting the elements of $UC(T)$ as the winners of T has been independently suggested by P. Fishburn in 1977 (Fishburn 1977) and by N. Miller in 1980 (Miller 1980).

For the tournament of Fig. 7, it is easy to see that z is covered by y , while x , y , and t are uncovered: here, x , y , and t are the three winners according to this solution.

The definition of a covered vertex can easily be extended to weighted tournaments, though the

possibility of weights equal to 0 makes the situation more difficult than expected (see Charon and Hudry (2006) for details).

Minimal Covering Set: Dutta's Solution

The uncovered set can also be refined, for example, by the minimal covering set proposed by B. Dutta (1988). Consider a tournament T defined on X . We say that a subset Y of X is a covering set of T if we have the following property: $\forall x \in X - Y, x \notin UC(Y \cup \{x\})$. For instance, $UC(T)$ is a covering set of T . Let $\mathcal{E}(T)$ denote the set of the covering sets of T . Then there exists a minimal element of $\mathcal{E}(T)$ with respect to inclusion; this minimal element is called the minimal covering set $MC(T)$ of T .

For the tournament of Fig. 7, the minimal covering set contains x , y , and t , which are the three winners according to this solution, here as for the uncovered set. (But there exist tournaments for which the general inclusion $MC(T) \subseteq UC(T)$ is strict.)

Maximal Transitive Subtournaments: Banks's Solution

Among other tournament solutions, one designed by J. Banks in 1985 (Banks 1985) is of interest. When the tournament being considered, T , is transitive (i.e., T is a linear order), there exists a unique winner who is selected as the winner of T by the usual tournament solutions. If that is not the case, we may consider the transitive subtournaments of T which are maximal with respect to inclusion and then select the winner of each of them as the winners of T . This defines the Banks's solution (Banks 1985): a Banks winner of T is the winner of any maximal (with respect to inclusion) transitive subtournament of T .

If we consider the majority tournament of Example 2 (see Fig. 7), three Banks winners exist: x (because of the maximal transitive subtournament $x > y > z$), y (because $y > z > t$), and t (because $t > x$).

Tournament Equilibrium Set: T. Schwartz's Solution

The solution that we deal with in this subsection was designed by T. Schwartz (1990) and is called

the tournament equilibrium set. To define it, we need extra definitions. Let G be a directed graph. The top set $TS(G)$ of G is defined as the union of the strongly connected components of G with no in-coming arcs (if G is a tournament, then $TS(G)$ is equal to $TC(G)$).

Let Sol be a tournament solution and T be a tournament. When Sol is applied to T , we define the contestation graph $G(Sol, T)$ associated with Sol and T as follows: The vertex set of $G(Sol, T)$ is X ; there is an arc (x, y) from x to y in $G(Sol, T)$ if and only if (x, y) is an arc of T and x is a winner according to Sol when Sol is applied to the subtournament of T induced by the predecessors of y in T . In other words, the arcs (x, y) of $G(Sol, T)$ describe the following situation: if y is considered as a possible winner, then x will contest the election of y because x beats y in T and because x is a winner among the candidates who beat y . We may notice that $G(Sol, T)$ is a subgraph of T . By considering the top set $TS[G(Sol, T)]$ of $G(Sol, T)$, we get a new tournament solution. T. Schwartz proved in Schwartz (1990) that there exists a unique tournament solution that he called the tournament equilibrium set TEQ , which is a fixed point with respect to this process: $\forall T, TEQ(T) = TS[G(TEQ, T)]$.

For the majority tournament of Example 2 (see Fig. 7), the tournament equilibrium set contains x , y , and t .

Dodgson's Procedure

In the procedure proposed in 1876 by C. L. Dodgson (or Lewis Carroll) (Dodgson 1876), each voter ranks all the candidates into a linear order. If a Condorcet winner exists, he or she is also the Dodgson winner. Otherwise, Dodgson's procedure consists of choosing as winners all the candidates who are "closest" to being Condorcet winners: for each candidate $x \in X$, let $D(x)$ be the Dodgson score of x , defined as the minimum number of swaps between consecutive candidates in the preferences of the voters such that x becomes a Condorcet winner; a Dodgson winner is any candidate x^* minimizing $D: D(x^*) = \min_{x \in X} D(x)$.

Specifically, consider a profile $\Pi = (L_1, L_2, \dots, L_m)$ of m linear orders. Let i be an

index between 1 and m and let y and z be two consecutive candidates in L_i : yL_iz , and there is no candidate t with yL_it and tL_iz simultaneously. Then define a linear order L'_i obtained from L_i by swapping y and z in L_i : xL'_iy , and for any other ordered pair of candidates $(t, v) \neq (y, z)$, we have tL'_iv if and only if we had tL_iv . By substituting L'_i to L_i , we obtain a new profile. By repeating such swaps as many times as necessary, we may generate all the possible profiles with m linear orders. Among them, some admit x as a Condorcet winner. The Dodgson score $D(x)$ of a candidate x is the minimum number of swaps that must be applied from Π in order to obtain a profile with x as a Condorcet winner.

Let us illustrate these definitions on Example 2. First, let us compute the Dodgson scores of the four candidates. If we want x to become a Condorcet winner, it is necessary that x defeats t , which is not currently the case. Since x is never just after t in the preferences of the voters, one swap is not enough to reach this aim. On the other hand, two swaps are enough, for instance, by swapping first y and x in the preference of the last voter and then by swapping x and t in this new preference. Thus, $D(x) = 2$. One swap is sufficient (and necessary) to make y a Condorcet winner, for instance, by swapping y and x in the preference of one of the first three voters: $D(y) = 1$. The computation of $D(z)$ is a little more difficult. To become a Condorcet winner, z must defeat y . Since there are only three voters who prefer z to y against six who prefer y to z , it is necessary to perform at least two swaps so that a new majority of voters will prefer z to y . But, if such swaps exist, they do not help to alleviate the difficulty in helping z to defeat x . For z to defeat x necessitates at least one extra swap. Hence, the inequality $D(z) \geq 3$. On the other hand, swapping x and z in the preference of the voter whose preference is $x > z > y > t$ and swapping y and z in the preferences of two of the first three voters make z a Condorcet winner: $D(z) = 3$. The computation of $D(t)$ is similar to that of $D(z)$, and we also find $D(t) = 3$. These values show that there is only one Dodgson winner (who is not a Condorcet-Kemeny winner): y .

Young's Procedure

Like Dodgson's procedure, H. P. Young's procedure (Young 1977) is based on altered profiles. But, instead of performing the fewest possible number of swaps of consecutive candidates in the voters' preferences, we now find the minimum number of voters that must be removed in order for a Condorcet winner to emerge. More precisely, for each candidate x , we define the Young score $Y(x)$ of x as the minimum number of voters whose simultaneous removals allow x to become a Condorcet winner. Any candidate with a minimum Young score is a Young winner.

For Example 2, we may verify that the Young scores of the four candidates are $Y(x) = 2$, $Y(y) = 2$, $Y(z) = 4$, and $Y(t) = 4$. For instance, for z to defeat x and y requires at least four removals for y . But these removals must keep z defeating t . It is easy to see that removing two of the first three voters, one of the two voters sharing the same preferences and the voter for whom z is in the last position, suffices to make z a Condorcet winner of the new election.

Approval Voting Procedure, Majority-Choice Approval Procedure, and Variants

The approval voting procedure was popularized by Steven J. Brams and Peter C. Fishburn in 1978 (Brams and Fishburn 1978) and 1983 (Brams and Fishburn 1983). But, according to Colomer and McLean (1998), Cox (1987), and Lines (1986), it was already in use in the thirteenth century in Venice and in papal elections and then in elections in England during the nineteenth century, among other places. Its name seems to come from R. J. Weber in 1976 (see Weber 1995); several other persons seem to have found this procedure independently during the late 1960s and early 1970s. The approval voting procedure consists of giving to the voters the possibility of voting for several candidates simultaneously. In other words, instead of appointing their preferred candidate, the voters are invited to answer the question: "who are the candidates for whom you would like to vote?" Each voter then gives 1 point to the candidates with whom he or she agrees. The candidate with the greatest number of points is the winner. Notice that, if each voter actually chooses only one candi-

date, this voting procedure is the same as the plurality rule. In this respect, the approval voting procedure can be seen as a generalization of the plurality rule.

More formally, this procedure assumes that the preferences of the voters are complete preorders with only two classes (one class for the approved candidates and one for the disapproved ones). These preorders must then be aggregated into a collective complete preorder also with two classes, one of them with only one element (the winner) and the other class with all the other elements (the losers). Extension for electing several candidates simultaneously is immediate.

Variants (sometimes known as range voting, ratings summation, average voting, cardinal ratings, 0–99 voting, the score system, or the point system) can be based on the same idea. For instance, each voter has a maximum number of points, and he or she can share them out among the candidates as he or she pleases, with or without a constraint on the maximum number of points per candidate. We can also add several rounds and thus define new procedures. The majority-choice approval procedure (MCA) designed by F. Simmons in 2002 can also be seen as a variant of approval voting. In this system, each voter has three possibilities for rating each candidate: “favored,” “accepted,” or “disapproved.” If a candidate is ranked “favored” by an absolute majority of the voters, then any candidate marked “favored” by a maximum number of voters is a winner. Otherwise, the winner is any candidate with the largest number of “favored” or “accepted” marks. It is sometimes required that this number be at least $(m + 1)/2$; otherwise, no one will be elected. Ties can be broken using the number of “favored” marks. We may, of course, obtain new variants by increasing the number of levels.

Bucklin’s Procedure

The procedure proposed by James W. Bucklin in the early twentieth century is also called the Grand Junction system (Grand Junction is a city in Colorado, where Bucklin’s procedure was applied from 1909 to 1922). In this procedure, we first count, for each candidate x , the number of times that x is ranked first, as in the plurality rule. If a

candidate has at least the absolute majority $(m + 1)/2$, he or she is the winner. Otherwise, second choices are added to the first choices. Once again, if a candidate has the absolute majority $(m + 1)/2$, he or she becomes a winner. Otherwise, we consider third choices and so on, until at least one candidate obtains the absolute majority: then he or she becomes a winner. Since after the first round there are more votes than voters, Bucklin’s procedure is sometimes considered as a variant of approval voting procedure, but a main difference is that Bucklin’s procedure may require several rounds. It can also be considered as a variant of scoring procedures, since it is the same as iteratively applying a scoring procedure with successively $(1, 0, 0, \dots, 0)$, $(1, 1, 0, 0, \dots, 0)$, $(1, 1, 1, 0, \dots, 0)$, $\dots$ as the score vectors, until a candidate obtains the majority.

Once again, consider Example 2. In the first round, x gets 4 points; y , 2 points; z , 1 point; t , 2 points. As there is no candidate with at least 5 points, we consider the second choices. Then we obtain: x keeps 4 points; y obtains 6 points; z , 3 points; t , 5 points. As y and t obtain at least 5 points, the process stops here and y and t are the Bucklin winners.

A variant of Bucklin’s procedure would be to consider, at the last iteration, the candidates with a maximum number of points as the winners (only y for Example 2).

Complexity Results

After a reminder of the complexity classes that are useful for our purposes, this section examines the complexity of the voting procedures described in the previous section (see also Faliszewski et al. 2009a). Among those tournament solutions which can be applied to the majority tournament, we restrict ourselves to the solutions depicted above; the reader interested in the complexity of the other common tournament solutions will find some answers in Hudry (2009).

Main Complexity Classes

Some complexity classes $\mathcal{Z}$ are well known: $\mathcal{P}$, $\mathcal{NP}$, $\text{co-}\mathcal{NP}$, and the sets of polynomial problems, of

NP-complete problems, and of NP-hard problems. . . . Let us recall some other notations (for references upon the theory of complexity, see, for instance Ausiello et al. 2003; Garey and Johnson 1979; Hemaspaandra 2000; Johnson 1990); see also Aaronson and Kuperberg (2013) for a list of about 500 complexity classes. Like D. S. Johnson in Johnson (1990), we will distinguish between decision problems and other types of problems. Let n denote the size of the data. The class L^{NP} , or $P^{NP[\log]}$, $P^{NP[\log n]}$, Θ_2^P , or also $P_{||}^{NP}$, contains those decision problems which can be solved by a deterministic Turing machine with an oracle in the following manner: the oracle can solve an appropriate NP-complete problem in unit time, and the number of consultations of the oracle is upper bounded by $\log(n)$. The class P^{NP} or Δ_2^P (or sometimes simply Δ_2) or also $P^{NP[n^{o(1)}]} = \bigcup_{k \geq 0} P^{NP[n^k]}$ is defined similarly but with a polynomial of n instead of $\log(n)$. Notice the inclusions: $NP \cup co - NP \subseteq L^{NP} \subseteq P^{NP}$. For problems which are not decision problems (i.e., optimization problems in which we look for the optimal value of a given function f , or search problems in which we look for an optimal solution of f , or enumeration problems in which we look for all the optimal solutions of f ; these problems are sometimes called function problems), these classes are extended with an “F” in front of their names (see Johnson 1990); for instance, we thus obtain the classes $FL^{NP}(= F\Theta_2^P \dots)$ and $FP^{NP}(= F\Delta_2^P \dots)$. The usual definition of “completeness” is extended to these classes in a natural way. Though the notation Δ_2^P or Θ_2^P is sometimes more common in complexity theory (especially when dealing with the polynomial hierarchy), we shall keep the pseudonyms P^{NP} and L^{NP} (as well as FP^{NP} and FL^{NP}), as perhaps being more informative.

Similarly, there exist refinements of P . In particular, the class L denotes the subset of P containing the (decision) problems which can be solved by an algorithm using only logarithmic space (in a deterministic Turing machine), the input itself not being counted as part of memory. The classes AC^0 and TC^0 are more technical. The first one, AC^0 , consists of all the problems solvable by uniform constant-depth Boolean circuits

with unbounded fan-in and a polynomial number of gates. The second one, TC^0 , consists of all the problems solvable by polynomial-size, bounded-depth, and unbounded fan-in Boolean circuits augmented by so-called “threshold” gates, i.e., unbounded fan-in gates that output “1” if and only if more than half their inputs are nonzero (see Johnson (1990) and the references therein for details). A problem of TC^0 is said to be TC^0 -complete if it is complete under AC^0 Turing reductions. Notice the inclusions $AC^0 \subset TC^0 \subseteq L \subseteq P$.

Complexity Results for the Usual Voting Procedures

We can now examine complexity results of the voting procedures described above. When dealing with linear orders, we assume that we can have direct access to the ordered list of the candidates, especially to their winners (see above). Remember that n denotes the number of candidates and m the number of voters.

Theorem 1 The following procedures are polynomial.

- The plurality rule (one-round procedure) is in $O(n + m)$.
- The plurality rule with runoff (two-round procedure) is in $O(n + m)$.
- The preferential voting procedure (STV) is in $O(nm + n^2)$.
- Borda’s procedure is in $O(nm)$.
- If the preferences of the voters are known through a profile of complete preorders with two classes, the approval voting procedure (and its variant suggested by F. Simmons) is in $O(nm)$.
- Bucklin’s procedure is in $O(nm)$.

More generally, the previous polynomial results can usually be extended to procedures based on score vectors, often with a complexity of $O(nm)$.

The previous results are rather easy to obtain. For instance, for the plurality rule, an efficient

way to determine the winners consists of computing an n -vector V which provides, for each candidate x , the number of voters who prefer x . As the preferred candidate of each voter is assumed to be accessible directly, computing V requires $O(n + m)$ operations (or fewer if we assume that voters sharing similar preferences are gathered, see subsection “[Definitions, Notation and Partially Ordered Sets Used to Model Preferences](#)”). More precisely, we first initialize V to 0 in $O(n)$. Then, for each voter, we consider his or her preferred candidate x , and we increment $V(x)$ by 1 in $O(1)$ for each voter and thus in $O(m)$ for all the voters. By scanning V in $O(n)$, we determine the maximum value contained in V and, once again by scanning V in $O(n)$, the winners: they are the candidates x whose values $V(x)$ are maximum. The whole process requires $O(n + m)$ operations.

When it succeeds in finding an order (or at least a Condorcet winner), Condorcet’s procedure is also polynomial. This is also the case for Simpson’s procedure, since computing the pairwise comparison coefficients m_{xy} can obviously be done in polynomial time. The same happens for Tideman’s ranked pairs procedure, since the detection of a circuit in a graph can be done in a time proportional to the number of arcs of this graph, i.e., here in $O(n^2)$ (see Cormen et al. 1990), and for the prudent order procedure or the minimax procedure.

Theorem 2 When there exists a Condorcet winner, Condorcet’s procedure is polynomial, in $O(n^2m)$. The prudent orders procedure, the minimax procedure, Simpson’s procedure (maximin procedure), and Tideman’s procedure (ranked pairs procedure) are also polynomial.

When there is no Condorcet winner, the situation is more difficult to manage. In this case, we may pay attention to the problems called $P_m(\mathcal{Y}, \mathcal{Z})$ above. The problems $P_m(\mathcal{Y}, \mathcal{T})$ and $P_m(\mathcal{Y}, \mathcal{R})$, i.e., the aggregation of m preferences into a tournament or into a binary relation in which no special property is required, are polynomial for any m and any set $\mathcal{Y}$ as specified by Theorem 3.

Theorem 3 Let m be any integer with $m \geq 1$ and let $\mathcal{Y}$ be any subset of $\mathcal{R}$. Consider a profile $\Pi \in \mathcal{Y}^m$. Then $P_m(\mathcal{Y}, \mathcal{R})$ and $P_m(\mathcal{Y}, \mathcal{T})$ are polynomial.

For the Condorcet-Kemeny problem, and more generally for the problems $P_m(\mathcal{Y}, \mathcal{Z})$ with $\mathcal{Y}$ belonging to $\{\mathcal{A}, \mathcal{C}, \mathcal{L}, \mathcal{O}, \mathcal{P}, \mathcal{R}, \mathcal{T}\}$ and $\mathcal{Z}$ to $\{\mathcal{A}, \mathcal{C}, \mathcal{L}, \mathcal{O}, \mathcal{P}\}$ and where m is a positive integer, most are NP-hard when m is large enough (see Alon 2006; Barthélemy et al. 1989; Bartholdi et al. 1989a; Charbit et al. 2007; Conitzer 2006; Dwork et al. 2001; Hemaspaandra et al. 2005; Hudry 1989, 2008, 2010, 2012, 2013b; Wakabayashi 1986, 1998); see also Hudry (2008) for the NP-hardness of similar problems $P_m(\mathcal{Y}, \mathcal{Z})$ when extended to other posets, including interval orders, interval relations, semiorders, weak orders, and quasi-orders and Hudry (2013a) for results dealing with other kinds of remoteness.

Theorem 4 For m large enough, we have:

- For any set $\mathcal{Y}$ containing $\mathcal{L}$ (this is the case in particular for $\mathcal{Y}$ belonging to $\{\mathcal{A}, \mathcal{C}, \mathcal{L}, \mathcal{O}, \mathcal{P}, \mathcal{R}, \mathcal{T}\}$ or to unions or intersections of such sets), $P_m(\mathcal{Y}, \mathcal{Z})$ is NP-hard, and the decision problem associated with $P_m(\mathcal{Y}, \mathcal{Z})$ is NP-complete for $\mathcal{Z} \in \{\mathcal{A}, \mathcal{C}, \mathcal{L}\}$.
- $P_m(\mathcal{R}, \mathcal{Z})$ is also NP-hard, and the decision problem associated with $P_m(\mathcal{R}, \mathcal{Z})$ is also NP-complete for $\mathcal{Z} \in \{\mathcal{O}, \mathcal{P}\}$.
- The complexity of $P_m(\mathcal{Y}, \mathcal{Z})$ is unknown for $\mathcal{Y} \in \{\mathcal{A}, \mathcal{C}, \mathcal{L}, \mathcal{O}, \mathcal{P}, \mathcal{T}\}$ and for $\mathcal{Z} \in \{\mathcal{O}, \mathcal{P}\}$, but $P_m(\mathcal{Y}, \mathcal{O})$ and $P_m(\mathcal{Y}, \mathcal{P})$ have the same complexity.

The minimum value of m for which $P_m(\mathcal{Y}, \mathcal{Z})$ is NP-hard in Theorem 4 is usually unknown, and moreover, the parity of m plays a role. Table 2 gives the ranges of lower bounds of m from which $P_m(\mathcal{Y}, \mathcal{Z})$ is known to be NP-hard; for lower values of m , when not trivial, the complexity of $P_m(\mathcal{Y}, \mathcal{Z})$ is usually unknown. (This is the case, for instance, for $P_1(\mathcal{R}, \mathcal{C})$, $P_1(\mathcal{R}, \mathcal{O})$, or $P_1(\mathcal{R}, \mathcal{P})$; notice that $P_2(\mathcal{L}, \mathcal{L})$ is polynomial.) A question mark (?) means that the complexity of the problem is still unknown.

Voting Procedures, Complexity of, Table 2 Lower bounds of m from which $P_m(\mathcal{Y}, \mathcal{Z})$ is known to be NP-hard

	$\mathcal{L} \subseteq \mathcal{Y}$	$\mathcal{L} \subseteq \mathcal{Y}$	$\mathcal{Y} = \mathcal{T}$	$\mathcal{Y} = \mathcal{T}$	$\mathcal{Y} = \mathcal{R}$	$\mathcal{Y} = \mathcal{R}$
Median relation ($\mathcal{Z}$)	m odd	m even	m odd	m even	m odd	m even
Acyclic relation ($\mathcal{A}$)	$\theta(n^2)$	4	1	2	1	2
Complete preorder ($\mathcal{C}$)	$\theta(n^2)$	4	1	2	1	2
Linear order ($\mathcal{L}$)	$\theta(n^2)$	4	1	2	1	2
Partial order ($\mathcal{O}$)	?	?	?	?	3	2
Preorder ($\mathcal{P}$)	?	?	?	?	3	2

So the Condorcet-Kemeny problem $P_m(\mathcal{L}, \mathcal{L})$ is NP-hard for m odd and large enough for m even with $m \geq 4$ (and, as noted above, is polynomial for $m = 2$). Similarly, $P_m(\mathcal{C}, \mathcal{C})$ (i.e., the problem considered by J. G. Kemeny in Kemeny (1959): the aggregation of complete preorders into a median complete preorder) is NP-hard for m odd and large enough for m even with $m \geq 4$ (see Hudry (2012) for the proof and for other results dealing with median complete preorders or with median weak orders). More specific results deal with $P_m(\mathcal{L}, \mathcal{L})$ (see Bartholdi et al. 1989a; Hemaspaandra et al. 2007; Hudry 2013b). To state them, let us define, for each element x of X and for any given profile Π , the Condorcet-Kemeny score $K(x)$ of x (with respect to Π) as the minimum remoteness $\Delta(\Pi, L_x)$ between Π and the linear orders L_x with x as their winner. The Condorcet-Kemeny index of Π , $K(\Pi)$ is the minimum remoteness between Π and any linear order; thus, it is the minimum that we look for in $P_m(\mathcal{L}, \mathcal{L})$; it is also the minimum of the Condorcet-Kemeny scores over X .

Theorem 5 Let Π be a profile of m linear orders with m large enough.

- The following problems are NP-complete:
 - Given Π , a candidate $x \in X$ and an integer k , is $K(x)$ lower than or equal to k ?
 - Given an integer k , is $K(\Pi)$ lower than or equal to k ?
- The following problems are NP-hard:
 - Given Π and two candidates $x \in X$ and $y \in X$ with $x \neq y$, is $K(x)$ lower than or equal to $K(y)$ (in other words, is x “better” than y)? Moreover, if we assume that the

profiles Π are given by the ordered list of the m preferences of the voters (and not by a set of different linear orders with their multiplicities; see subsection “Definitions, Notation and Partially Ordered Sets Used to Model Preferences”), this problem is L^{NP} -complete.

- Given Π and a candidate $x \in X$, is x a Condorcet-Kemeny winner? Moreover, if we assume that the profiles Π are given by the ordered list of the m preferences of the voters, this problem is L^{NP} -complete.
 - Given Π and a candidate $x \in X$, is x the unique Condorcet-Kemeny winner? Moreover, if we assume that the profiles Π are given by the ordered list of the m preferences of the voters, this problem is L^{NP} -complete.
 - Given Π , determine a Condorcet-Kemeny winner of Π . Moreover, this problem belongs to FP^{NP} .
 - Given Π , determine all the Condorcet-Kemeny winners of Π . Moreover, this problem belongs to FP^{NP} .
 - Given Π , determine a Kemeny order of Π . Moreover, this problem belongs to FP^{NP} .
- The following problem belongs to the class co-NP:
 - Given Π and a linear order L defined on X , is L a median linear order of Π ?

Attention has also been paid to $P_1(\mathcal{T}, \mathcal{L})$, i.e., the approximation of a tournament (which can be the majority tournament of the election) by a linear order at minimum distance (Slater problem (Slater 1961)). The complexity of the Slater problem derives from a recent result dealing with the problem called the feedback arc set problem,

which is known to be NP-complete from work by R. Karp (1972) for general graphs. This problem consists of removing a minimum number of arcs in a directed graph in order to obtain a graph without any circuits. Recent results (Alon 2006; Charbit et al. 2007; Conitzer 2006) show that this problem remains NP-complete even when restricted to tournaments. For a tournament T , removing a minimum number of arcs to obtain a graph without any circuits is the same as reversing a minimum number of arcs to make T transitive, i.e., a linear order (see Charon and Hudry (2007) and Hudry (2010)). From this, we may prove the following theorem (see Hudry (2010) for details):

Theorem 6 For any tournament T , we have the following results:

- The computation of the Slater index $i(T)$ of T is NP-hard; this problem belongs to the class FP^{NP} ; the associated decision problem is NP-complete.
- The computation of a Slater winner of T is NP-hard; this problem belongs to the class FP^{NP} .
- Checking that a given vertex is a Slater winner of T is NP-hard; this problem belongs to the class L^{NP} .
- The computation of a Slater order of T is NP-hard; this problem belongs to the class FP^{NP} .
- The computation of all the Slater winners of T is NP-hard; this problem belongs to the class FP^{NP} .
- The computation of all the Slater orders of T is NP-hard.
- Checking that a given order is a Slater order is a problem which belongs to the class co-NP.

Similarly, the computation of a median complete preorder of a profile of m tournaments (i.e., with the previous notation, $P_m(T, C)$) is NP-hard for any m with m greater than or equal to 1 (see Hudry 2012).

Some of the previous results may be generalized to other definitions of remoteness, for instance, if the sum of the symmetric difference distances to the individual preferences is replaced by their maximum, or by the sum of their squares,

or by the sum of any of their positive powers, as specified below (see Hudry 2013a):

Theorem 7 Let Φ denote any remoteness defined for any profile Π such that, when m is equal to 1 with $\Pi = (R_1)$, then the minimization of $\Phi(\Pi, R)$ over the relations R belonging to $\mathcal{A}$ yields to the same optimal solutions as the minimization of $\delta(R_1, R)$ over the same set. Then, for any fixed m with $m \geq 1$, the aggregation of m binary relations or m tournaments obtained by minimizing Φ (the problems similar to $P_m(\mathcal{R}, \mathcal{Z})$ and $P_m(\mathcal{T}, \mathcal{Z})$ for $\mathcal{Z} \in \{\mathcal{A}, \mathcal{L}\}$ but with respect to Φ) is NP-hard, and the associated decision problems are NP-complete.

An important case exists for which $P_m(\mathcal{L}, \mathcal{L})$ becomes polynomial for any m : it is the one for which the voters' preferences are single-peaked linear orders. To define them, assume that we can order the candidates on a line, from left to right, and assume that this linear order Ω does not depend on the voters (from a practical point of view, Ω is not always easy to define, even for political elections). For any voter α , let x_α denote the candidate preferred by α . The preference of α is said to be Ω -single-peaked if, for any candidates y and z with $x \neq y \neq z \neq x$ and located on the same side of Ω from x_α , y is preferred to z by α if and only if y is closer to x_α than z with respect to Ω . Let $\mathcal{U}_\Omega$ denote the set of Ω -single-peaked linear orders. D. Black (1958) showed that, for any order Ω , the aggregation of Ω -single-peaked linear orders is an Ω -single-peaked linear order (see also Conitzer 2007 for another study dealing with single-peaked preferences). Hence, the polynomiality of the aggregation of Ω -single-peaked linear orders into an Ω -single-peaked linear order:

Theorem 8 For any linear order Ω defined on X and any positive integer m , for any set $\mathcal{Z}$ containing $\mathcal{U}_\Omega$ as a subset (this is the case for the sets $\mathcal{A}, \mathcal{C}, \mathcal{L}, \mathcal{O}, \mathcal{P}, \mathcal{R}, \mathcal{T}$), $P_m(\mathcal{U}_\Omega, \mathcal{Z})$ is polynomial. More precisely, $P_m(\mathcal{U}_\Omega, \mathcal{Z})$ can be solved in $O(n^2 m)$.

The NP-hardness of Slater's problems shows that tournament solutions can be difficult to compute. This obviously depends on the solutions

considered. For instance, Copeland's procedure is polynomial, as specified by the next theorem.

Theorem 9 Copeland's procedure is polynomial, in $O(n^2m)$.

If we assume that the tournament related to Copeland's procedure is already computed, and if we do not take the memory space necessary to code this tournament into account, it is easy to see that the memory space necessary to compute the maximum of the Copeland scores and then to decide whether a given vertex is a Copeland winner can be bounded by a constant. So deciding whether a given vertex is a Copeland winner is a problem belonging to class L. In fact, a stronger result is shown by F. Brandt, F. Fischer, and P. Harrenstein in Brandt et al. (2006):

Theorem 10 Checking that a given vertex is a Copeland winner is a TC^0 -complete problem.

Similar results can be stated for Smith's tournament solution (top cycle):

Theorem 11 Let T be a tournament.

- The computation of the Smith winners of T (the elements of the top cycle of T) can be done in $O(n^2)$ (or even in linear time with respect to the cardinality of the top cycle if the sorted scores are known).
- Checking that a given vertex is a Smith winner is a TC^0 -complete problem (Brandt et al. 2006).

For the other tournament solutions depicted above, we have the results stated in Theorem 12 (see Hudry 2009 for details and for results upon other tournament solutions). This theorem shows that checking whether a given vertex of a given tournament is a Banks winner is NP-complete (Woeginger 2003) (and Brandt et al. 2008 for an alternative proof), while computing a Banks winner is polynomial (Hudry 2004). Of course, when such a Banks winner is computed, we do not choose the winner that we compute among the set of Banks winners (and so there is no

contradiction between the two results if P and NP are different). The polynomiality of the minimal covering set solution (through n resolutions of a linear programming problem, which can be done in polynomial time using L. Khachiyan's algorithm (Khachiyan 1979)) and the NP-hardness of the tournament equilibrium set solution are new results, respectively, due to F. Brandt and F. Fischer (2008) and to Brandt et al. (2008).

Theorem 12 Let T be a tournament.

- Computing the uncovered elements of T can be done within the same complexity as the multiplication of two $(n \times n)$ matrices and so can be done in $O(n^{2.38})$ operations.
- Computing the elements of the minimal covering set of T can be done in polynomial time with respect to n .
- The following problem is NP-complete: given a tournament T and a vertex x of T , is x a Banks winner of T ?
- Computing a Banks winner is polynomial and, more precisely, can be done in $O(n^2)$ operations.
- Computing all the Banks winners of T is an NP-hard problem. More precisely, it is a problem belonging to the class FP^{NP} .
- The following decision problem is NP-hard (but is not known to be inside NP): given a vertex x of T , does x belong to $TEQ(T)$?

J. J. Bartholdi III, C. A. Tovey, and A. Trick stated in Bartholdi et al. (1989a) that Dodgson's procedure is NP-hard. More precisely, they proved the NP-completeness of the problem of Theorem 13 and the NP-hardness of the first two problems of Theorem 14. Hemaspaandra et al. (1997) sharpened their results (Theorem 14), assuming that Π is given by the ordered list of the m preferences of the voters (and not by a set of different linear orders with their multiplicities).

Theorem 13 The following problem is NP-complete: given a profile Π of linear orders defined on X , a candidate $x \in X$, and an integer k , is the Dodgson score $D(x)$ of x less than or equal to k ?

Theorem 14 The following problems are NP-hard and, more precisely, L^{NP} -complete if the considered profiles Π are given by the ordered lists of the m preferences of the voters.

- Given a profile Π of linear orders defined on X and a candidate $x \in X$, is x a Dodgson winner?
- Given a profile Π of linear orders defined on X and a candidate $x \in X$, is x the unique Dodgson winner?
- Given a profile Π of linear orders defined on X and two candidates $x \in X$ and $y \in X$ with $x \neq y$, is $D(x)$ less than or equal to $D(y)$ (in other words, is x “better” than y)?

Notice that a variant of Dodgson’s procedure, called homogeneous Dodgson’s procedure (see Fishburn 1977; Vazirani 2003), has been shown to be polynomial by Rothe et al. (2003). Consider that each voter is replicated p times, each copy having the same preference. A procedure is said to be homogeneous if such a replication does not change the winners. Dodgson’s procedure and Young’s procedure are not homogeneous (Fishburn 1977). Then, instead of the Dodgson score defined above, we may consider, for each candidate x , the limit, when p tends to infinity, of the ratio between, on the one hand, the Dodgson score of x after the replication of each voter p times and, on the other hand, p . J. Rothe, H. Spakowski, and J. Vogel provide in Rothe et al. (2003) a linear program for computing such a limit for any candidate, hence the polynomiality of the homogeneous version of Dodgson’s procedure.

They give, in the same paper (Rothe et al. 2003), the complexity of Young’s procedure.

Theorem 15 The following problems are NP-hard and, more precisely, L^{NP} -complete.

- Given a profile Π of linear orders defined on X and a candidate $x \in X$, is x a Young winner?
- Given a profile Π of linear orders defined on X and two candidates $x \in X$ and $y \in X$ with $x \neq y$, is $Y(x)$ less than or equal to $Y(y)$ (in other words, is x “better” than y)?

Notice that all these problems become obviously polynomial if we assume that the number of candidates is upper bounded by a constant.

Further Directions

Complexity is a prominent feature of voting procedures. NP-hardness may involve too much important CPU time to solve a given instance if the number of candidates is not small. In this respect, polynomial procedures can be preferable to NP-hard ones. In practice, the number of candidates is not always too large, and the instance considered can be tractable. This also depends on the efficiency of the applied algorithms (see Barthélemy and Monjardet 1981; Charon and Hudry 2006, 2007; Charon et al. 1997; Conitzer 2006; Homan and Hemaspaandra 2006; Hudry 1989; Jünger 1985; LeGrand et al. 2006; Reinelt 1985; Rothe and Spakowski 2006; Wakabayashi 1986 for references on algorithms designed to solve some of the previous problems). Moreover, complexity is stated here in the worst case. A study about the average complexity (for appropriate probability distributions) remains to be done.

Other directions can be investigated.

For instance, the parameterized complexity (see Downey and Fellows (1999) for a global presentation of this field and Fischer et al. (2013) for results on this topic in the context of voting theory) of the feedback arc set problem applied to tournaments has also been studied. V. Raman and S. Saurabh (2006) (see also Dom et al. (2006), where the case of bipartite tournaments is also considered) show that the feedback arc set problem for weighted or unweighted tournaments is fixed-parameter tractable (FPT) by providing appropriate algorithms. (Remember that a problem is said to be FPT if there exist an arbitrary function f and a polynomial function Q such that for any instance (I, k) where I is an instance of the non-parameterized version of the problem and k is the considered complexity parameter, (I, k) can be solved within a CPU time upper bounded by $f(k)Q(|I|)$, where $|I|$ denotes the size of I . In particular, if k is upper bounded by a constant, then the problem becomes polynomial.) V. Raman

and S. Saurabh give several algorithms to solve the parameterized version of the feedback arc set problem applied to tournaments. The best complexity of their algorithms for this problem is $O(2.415^k n^\omega)$, where ω denotes the exponent of the running time of the best matrix multiplication algorithm (for instance, in the method designed by D. Coppersmith and S. Winograd 1987, ω is about 2.376). Parameterized complexity appears also in Christian et al. (2006), McCabe-Dansted (2006), and McCabe-Dansted et al. (2006) for Condorcet-Kemeny's, Dodgson's, and Young's procedures.

Another direction deals with algorithms with approximation guarantees (see Ausiello et al. (2003) and Vazirani (2003) for a presentation of this field and Charon and Hudry (2007) and Fischer et al. (2013) for results on this topic in the context of voting theory). For instance, while the feedback arc set problem is APX-hard in general (Even et al. 1998) (remember that, from a practical point of view, this implies that we do not know polynomial-time approximation scheme – PTAS – to solve this problem), Ailon et al. (2005) designed randomized 3-approximation and 2.5-approximation algorithms for the feedback arc set problem when restricted to unweighted tournaments. The 3-approximation algorithm has been derandomized by A. van Zuylen (2005), still for unweighted tournaments and with an approximation ratio equal to 3. These methods may be adapted to weighted tournaments with an approximation ratio equal to 5 (see Ailon et al. 2005; Coppersmith et al. 2006). N. Alon (2006) (see also Ailon and Alon 2007) shows that, for any fixed $\varepsilon > 0$, it is NP-hard to approximate the minimum size of a feedback arc set for a tournament on n vertices up to an additive error of $n^{2-\varepsilon}$ (but approximating it up to an additive error of εn^2 can be done in polynomial time). Ranking the vertices of a tournament according to their out-degrees (Copeland's procedure) for unweighted tournaments, or according to the sum of the weights of the arcs leaving them minus the sum of the weights of the arcs entering them (Borda's procedure) for weighted tournaments, also provides a method with approximation ratio equal to

5, irrespective of how ties are broken (see Coppersmith et al. 2006; Tideman 1987).

Probabilistic algorithms (see Alon and Spencer (2000) for a global presentation of these methods) have also been applied to the feedback arc set problem (or rather to the search for a maximum subdigraph without circuits, which is the same for tournaments), for unweighted (Czygrinow et al. 1999; de la Vega 1983; Poljak and Turzík 1986; Poljak et al. 1988; Spencer 1971, 1978, 1987) or weighted (Czygrinow et al. 1999) tournaments.

If NP-hardness is usually considered as a drawback (because the CPU time necessary to solve the instances becomes quickly prohibitive, because the algorithms may be difficult to explain to the voters, and so on), it can also be an asset with respect to manipulation, bribery, or other attempts to control the election (by adding or deleting candidates or voters). This is an emerging but already flourishing topic, which becomes a subject on its own (see Bartholdi and Orlin 1991; Bartholdi et al. 1989b, 1992; Chamberlin 1985; Conitzer and Sandholm 2002a, b, 2003, 2006; Conitzer et al. 2003; Elkin and Lipmaa 2006; Faliszewski et al. 2006, 2009a, b; Gibbard 1973; Hemaspaandra and Hemaspaandra 2007; Hemaspaandra et al. 2006, 2007; LeGrand et al. 2006; Maus et al. 2006; Mitlöhner et al. 2006; Moulin 1980, 1985; Procaccia and Rosenschein 2006; Procaccia et al. 2006; Saari 1990; Satterthwaite 1975; Smith 1999; Taylor 2005). A voting procedure is said to be manipulable when a voter, who knows how the other voters vote, has the opportunity to benefit by strategic or tactical voting, i.e., when the voter supports a candidate other than his or her sincerely preferred candidate in order to prevent an undesirable outcome (the main difference between manipulation and bribery is that for manipulability, the manipulators are known as a part of the instance, and for bribery, the number of corrupted voters is bounded, but these manipulators are not given). It is known from the theorem by A. Gibbard (1973) and M. Satterthwaite (1975) that, for at least three candidates, any voting procedure without a dictator is manipulable. Because of this result, manipulation cannot be precluded in any reasonable voting procedure on at least three candidates. As suggested in Bartholdi

and Orlin (1991) and in Bartholdi et al. (1989b) (see also Faliszewski et al. 2009a), the situation is less bad if we adopt a procedure for which manipulation is NP-hard. Some procedures are easy to manipulate: this is the case, for instance, for the plurality rule, Borda's procedure, maximin procedure, or Copeland's procedure (Bartholdi et al. 1989b) (see also Faliszewski et al. 2009a for variants). Some others are NP-hard: this is the case for some variants of some procedures based on score vectors (Conitzer and Sandholm 2002b; Conitzer et al. 2003; Hemaspaandra and Hemaspaandra 2007) or for other procedures (Faliszewski et al. 2006; Hemaspaandra et al. 2007). But we must remember that NP-hardness refers to the complexity in the worst case; V. Conitzer and T. Sandholm show in Conitzer and Sandholm (2006) that, under certain assumptions, no voting procedure is hard to manipulate, on average.

Among further related issues, we can still mention complexity issues in voting on combinatorial domains (see, for instance, Lang 2004) or the complexity of determining if the output of a vote is determined even when some voters have not yet been elicited (see, for instance, Conitzer and Sandholm 2002a).

Computational complexity has become an important criterion in evaluating voting procedures, even if real elections are not necessarily the most difficult instances (e.g., because the number of candidates can be limited). Voting methods are increasingly used in multiagent systems. When autonomous software agents are voting over all sorts of issues, then large numbers of candidates may be more likely and also agents are possibly more likely to manipulate. Thus, complexity is now a property which must be taken into account along with axiomatic properties. In spite of Arrow's impossibility theorem, or maybe because of it, it is still necessary to design new voting procedures, to study them, and to compare them to existing procedures. In the mathematical tool box available to do so, there is now a place for computational complexity.

Acknowledgments I would like to thank Ulle Endriss, Jérôme Lang, and Bernard Monjardet for their help. Their comments were very useful to improve the text.

Bibliography

Primary Literature

- Aaronson S, Kuperberg G (2013) Complexity zoo. https://complexityzoo.uwaterloo.ca/Complexity_Zoo
- Ailon N, Alon N (2007) Hardness of fully dense problems. *Inf Comput* 205:117–1129
- Ailon N, Charikar M, Newman A (2005) Aggregating inconsistent information: ranking and clustering. *Proceedings of the 37th annual ACM symposium on theory of computing (STOC)*, pp 684–693
- Aizerman MA, Aleskerov FT (1995) *Theory of choice*. North Holland, Elsevier, Amsterdam
- Alon N (2006) Ranking tournaments. *SIAM J Discret Math* 20(1):137–142
- Alon N, Spencer J (2000) *The probabilistic method*, 2nd edn. Wiley, New York
- Arrow KJ (1963) *Social choice and individual values*, rev edn. Wiley, New York
- Arrow KJ, Raynaud H (1986) *Social choice and multi-criterion decision-making*. MIT Press, Cambridge
- Arrow KJ, Sen AK, Suzumura K (eds) (2002) *Handbook of social choice and welfare*, vol 1. North-Holland, Amsterdam
- Ausiello G, Crescenzi P, Gambosi G, Kann V, Marchetti-Spaccamela A, Protasi M (2003) *Complexity and approximation*, 2nd edn. Springer, Berlin
- Baldwin JM (1926) The technique of the Nanson preferential majority system of election. *Proc R Soc Vic* 39:42–52
- Bang-Jensen J, Gutin G (2001) *Digraphs: theory, algorithms, and applications*. Springer, Berlin
- Banks J (1985) Sophisticated voting outcomes and agenda control. *Soc Choice Welf* 2:295–306
- Barnett WA, Moulin H, Salles M, Schofield NJ (eds) (1995) *Social choice, welfare and ethics*. Cambridge University Press, New York
- Barthélemy J-P (1979) Caractérisations axiomatiques de la distance de la différence symétrique entre des relations binaires. *Math Sci Hum* 67:85–113
- Barthélemy J-P, Monjardet B (1981) The median procedure in cluster analysis and social choice theory. *Math Soc Sci* 1:235–267
- Barthélemy J-P, Guénoche A, Hudry O (1989) Median linear orders: heuristics and a branch and bound algorithm. *Eur J Oper Res* 41:313–325
- Bartholdi JJ III, Orlin J (1991) Single transferable vote resists strategic voting. *Soc Choice Welf* 8(4):341–354
- Bartholdi JJ III, Tovey CA, Trick MA (1989a) Voting schemes for which it can be difficult to tell who won the election. *Soc Choice Welf* 6:157–165
- Bartholdi JJ III, Tovey CA, Trick MA (1989b) The computational difficulty of manipulating an election. *Soc Choice Welf* 6:227–241
- Bartholdi JJ III, Tovey CA, Trick MA (1992) How hard is it to control an election? *Math Comput Model* 16(8/9):27–40
- Berge C (1985) *Graphs*. North-Holland, Amsterdam
- Black D (1958) *The theory of committees and elections*. Cambridge University Press, Cambridge

- Borda J-C (1784) *Mémoire sur les élections au scrutin*. Histoire de l'Académie Royale des Sciences pour 1781, Paris, pp 657–665. English translation: de Grazia A (1953) Mathematical derivation of an election system. *Isis* 44:42–51
- Brams SJ, Fishburn PC (1978) Approval voting. *Am Polit Sci Rev* 72(3):831–857
- Brams SJ, Fishburn PC (1983) *Approval voting*. Birkhauser, Boston
- Brams SJ, Fishburn PC (2002) Voting procedures. In: Arrow K, Sen A, Suzumura K (eds) *Handbook of social choice and welfare*, vol 1. Elsevier, Amsterdam, pp 175–236
- Brandt F, Fischer F (2008) Computing the minimal covering set. *Math Soc Sci* 58(2):254–268
- Brandt F, Fischer F, Harrenstein P (2006) The computational complexity of choice sets. In: Endriss U, Lang J (eds) *Proceedings of the conference computational social choice 2006*. University of Amsterdam, Amsterdam, pp 63–76
- Brandt F, Fischer F, Harrenstein P, Mair M (2008) A computational analysis of the tournament equilibrium set. In: Fox D, Gomes CP (eds) *Proceedings of AAAI*, pp 38–43
- Caritat MJAN, marquis de Condorcet (1785) *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, Paris
- Casparid N, Monjardet B, Leclerc B (2007) *Ensembles ordonnés finis: concepts, résultats et usages*. Springer, Berlin
- Chamberlin JR (1985) An investigation into the effective manipulability of four voting systems. *Behav Sci* 30:195–203
- Charbit P, Thomassé S, Yeo A (2007) The minimum feedback arc set problem is NP-hard for tournaments. *Comb Probab Comput* 16(1):1–4
- Charon I, Hudry O (2006) A branch and bound algorithm to solve the linear ordering problem for weighted tournaments. *Discret Appl Math* 154:2097–2116
- Charon I, Hudry O (2007) A survey on the linear ordering problem for weighted or unweighted tournaments. *4OR* 5(1):5–60
- Charon I, Guénoche A, Hudry O, Woïgard F (1997) New results on the computation of median orders. *Discret Math* 165–166:139–154
- Chevalerey Y, Endriss U, Lang J, Maudet N (2007) A short introduction to computational social choice. In: *Proceedings of the 33rd conference on current trends in theory and practice of computer science (SOFSEM-2007)*. Lecture notes in computer science, vol 4362. Springer, Berlin, pp 51–69
- Christian R, Fellows M, Rosamond F, Slinko A (2006) On complexity of lobbying in multiple referenda. In: *Proceedings of the first international workshop on computational social choice (COMSOC 2006)*. University of Amsterdam, pp 87–96
- Colomer JM, McLean I (1998) Electing popes: approval balloting and qualified-majority rule. *J Interdiscip Hist* 29(1):1–22
- Conitzer V (2006) Computing Slater rankings using similarities among candidates. In: *Proceedings of the 21st national conference on artificial intelligence, AAAI-06*, Boston, pp 613–619
- Conitzer V (2007) Eliciting single-peaked preferences using comparison queries. In: *Proceedings of the 6th international joint conference on autonomous agents and multi agent systems (AAMAS-07)*, Honolulu, pp 408–415
- Conitzer V, Sandholm T (2002a) Vote elicitation: complexity and strategy-proofness. In: *Proceedings of the national conference on artificial intelligence (AAAI)*, pp 392–397
- Conitzer V, Sandholm T (2002b) Complexity of manipulating elections with few candidates. In: *Proceedings of the 18th national conference on artificial intelligence (AAAI)*, pp 314–319
- Conitzer V, Sandholm T (2003) Universal voting protocol tweaks to make manipulation hard. In: *Proceedings of the 18th international joint conference on artificial intelligence (IJCAI-03)*, Acapulco, pp 781–788
- Conitzer V, Sandholm T (2006) Nonexistence of voting rules that are usually hard to manipulate. In: *Proceedings of the 21st national conference on artificial intelligence (AAAI-06)*, Boston, pp 627–634
- Conitzer V, Lang J, Sandholm T (2003) How many candidates are needed to make elections hard to manipulate? Theoretical aspects of rationality and knowledge (TARK), pp 201–214
- Copeland AH (1951) A “reasonable” social welfare function. Seminar on applications of mathematics to the social sciences. University of Michigan
- Coppersmith T, Winograd S (1987) Matrix multiplication via arithmetic progression. In: *Proceedings of 19th annual ACM symposium on theory of computing*, pp 1–6
- Coppersmith D, Fleischer L, Rudra A (2006) Ordering by weighted number of wins gives a good ranking for weighted tournaments. In: *Proceedings of the 17th annual ACM-SIAM symposium on discrete algorithms (SODA'06)*, pp 776–782
- Cormen T, Leiserson C, Rivest R (1990) *Introduction to algorithms*, 2nd edn. MIT Press, Cambridge, 2001
- Cox GW (1987) *The cabinet and the development of political parties in Victorian England*. Cambridge University Press, New York
- Czygrinow A, Poljak S, Rödl V (1999) Constructive quasi-Ramsey numbers and tournament ranking. *SIAM J Discret Math* 12(1):48–63
- Daunou PCF (1803) *Mémoire sur les élections au scrutin*. Baudouin, Paris, an XI
- de la Vega WF (1983) On the maximal cardinality of a consistent set of arcs in a random tournament. *J Comb Theor B* 35:328–332
- Debord B (1987a) Caractérisation des matrices de préférences nettes et méthodes d'agrégation associées. *Math Sci Hum* 97:5–17
- Debord B (1987b) *Axiomatisation de procédures d'agrégation de préférences*. Ph D thesis, Université scientifique technologique et médicale de Grenoble

- Dodgson CL (1873) A discussion of the various methods of procedure in conducting elections. Imprint by Gardner EB, Hall EP, Stacy JH. Printers to the University, Oxford. Reprinted In: Black D (1958) *The theory of committees and elections*. Cambridge University Press, Cambridge, pp 214–222
- Dodgson CL (1874) Suggestions as to the best method of taking votes, where more than two issues are to be voted on. Imprint by Hall EP, Stacy JH. Printers to the University, Oxford. Reprinted In: Black D (1958) *The theory of committees and elections*. Cambridge University Press, Cambridge, pp 222–224
- Dodgson CL (1876) A method of taking votes on more than two issues. Clarendon Press, Oxford. Reprint In: Black D (1958) *The theory of committees and elections*. Cambridge University Press, Cambridge, pp 224–234; and In: McLean I, Urken A (1995) *Classics of social choice*. University of Michigan Press, Ann Arbor
- Dom M, Guo J, Hüffner F, Niedermeier R, Truß A (2006) Fixed-parameter tractability results for feedback set problems in tournaments, vol 3998, *Lecture notes in computer science*. Springer, Berlin, pp 320–331
- Downey RG, Fellows MR (1999) *Parameterized complexity*. Springer, Berlin
- Dummett M (1984) *Voting procedures*. Clarendon, Oxford
- Dutta B (1988) Covering sets and a new Condorcet choice correspondence. *J Econ Theor* 44:63–80
- Dwork C, Kumar R, Naor M, Sivakumar D (2001) Rank aggregation methods for the Web. In: *Proceedings of the 10th international conference on World Wide Web (WWW10)*, Hong Kong, pp 613–622
- Elkin E, Lipmaa H (2006) Hybrid voting protocols and hardness of manipulation. In: Endriss U, Lang J (eds) *Proceedings of the first international workshop on computational social choice (COMSOC 2006)*. University of Amsterdam, pp 178–191
- Elster J, Hylland A (eds) (1986) *Foundations of social choice theory*. Cambridge University Press, New York
- Erdős P, Moser L (1964) On the representation of directed graphs as unions of orderings. *Magyar Tud Akad Mat Kutató Int Közl* 9:125–132
- Even G, Naor JS, Sudan M, Schieber B (1998) Approximating minimum feedback sets and multicuts in directed graphs. *Algorithmica* 20(2):151–174
- Fagin R, Kumar R, Mahdian M, Sivakumar D, Vee E (2005) Rank aggregation: an algorithmic perspective. Unpublished manuscript
- Faliszewski P, Hemaspaandra E, Hemaspaandra L (2006) The complexity of bribery in elections. In: Endriss U, Lang J (eds) *Proceedings of the first international workshop on computational social choice (COMSOC 2006)*. University of Amsterdam, pp 178–191
- Faliszewski P, Hemaspaandra E, Hemaspaandra L, Rothe J (2009a) A richer understanding of the complexity of election systems. In: Ravi S, Shukla S (eds) *Fundamental problems in computing: essays in honor of Professor Daniel J. Rosenkrantz*. Springer, Berlin, pp 375–406
- Faliszewski P, Hemaspaandra E, Hemaspaandra L, Rothe J (2009b) Llull and Copeland voting broadly resist bribery and control. *J AI Res* 35:275–341
- Fischer F, Hudry O, Niedermeier R (2013) Weighted tournament solutions. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia A (eds) *Handbook of computational social choice*. Cambridge University Press, Cambridge, to appear
- Fishburn PC (1973a) Interval representations for interval orders and semiorders. *J Math Psychol* 10:91–105
- Fishburn PC (1973b) *The theory of social choice*. Princeton University Press, Princeton
- Fishburn PC (1977) Condorcet social choice functions. *SIAM J Appl Math* 33:469–489
- Fishburn PC (1985) *Interval orders and interval graphs, a study of partially ordered sets*. Wiley, New York
- Garey MR, Johnson DS (1979) *Computers and intractability, a guide to the theory of NP-completeness*. Freeman, New York
- Gibbard A (1973) Manipulation of voting schemes. *Econometrica* 41:587–602
- Guilbaud GT (1952) Les théories de l'intérêt général et le problème logique de l'agrégation. *Économie Appl* 5(4):501–584; *Éléments de la théorie des jeux*, 1968. Dunod, Paris
- Hägle G, Pukelsheim F (2001) Llull's writings on electoral systems. *Stud Lulliana* 3:3–38
- Hemaspaandra L (2000) Complexity classes. In: Rosen KH (ed) *Handbook of discrete and combinatorial mathematics*. CRC Press, Boca Raton, pp 1085–1090
- Hemaspaandra E, Hemaspaandra L (2007) Dichotomy for voting systems. *J Comput Syst Sci* 73(1):73–83
- Hemaspaandra E, Hemaspaandra L, Rothe J (1997) Exact analysis of Dodgson elections: Lewis Carroll's 1876 voting system is complete for parallel access to NP. *J ACM* 44(6):806–825
- Hemaspaandra E, Spakowski H, Vogel J (2005) The complexity of Kemeny elections. *Theor Comput Sci* 349:382–391
- Hemaspaandra E, Hemaspaandra L, Rothe J (2006) Hybrid elections broaden complexity-theoretic resistance to control. In: *Proceedings of the first international workshop on computational social choice (COMSOC 2006)*, University of Amsterdam, pp 234–247; (2007) *Proceedings of the 20th international joint conference on artificial intelligence (IJCAI 2007)*. AAAI Press, pp 1308–1314
- Hemaspaandra E, Hemaspaandra L, Rothe J (2007) Anyone but him: the complexity of precluding an alternative. *Artif Intell* 171(5–6):255–285
- Homan C, Hemaspaandra L (2006) Guarantees for the success frequency of an algorithm for finding Dodgson-election winners. In: *Proceedings of the 31st international symposium on mathematical foundations of computer science. Lecture notes in computer science*, vol 4162. Springer, Berlin, pp 528–539
- Hudry O (1989) Recherche d'ordres médians: complexité, algorithmique et problèmes combinatoires. Ph D thesis, ENSI, Paris
- Hudry O (2004) A note on Banks winners. In: Woeginger GJ (ed) *Tournaments are difficult to recognize*. *Soc Choice Welf* 23:1–2
- Hudry O (2008) NP-hardness results on the aggregation of linear orders into median orders. *Ann Oper Res* 163(1):63–88

- Hudry O (2009) A survey on the complexity of tournament solutions. *Math Soc Sci* 57:292–303
- Hudry O (2010) On the complexity of Slater's problems. *Eur J Oper Res* 203:216–221
- Hudry O (2012) On the computation of median linear orders, of median complete preorders and of median weak orders. *Math Soc Sci* 64:2–10
- Hudry O (2013a) Complexity results for extensions of median orders to different types of remoteness. *Ann Oper Res*. doi:10.1007/s10479-013-1342-3 to appear
- Hudry O (2013b) Complexity of computing median linear orders and variants. *Electron Notes Discrete Math* 42:57
- Hudry O, Monjardet B (2010) Consensus theories. An oriented survey. *Math Soc Sci* 190:139–167
- Hudry O, Leclerc B, Monjardet B, Barthélemy J-P (2009) Metric and latticial medians. In: Bouyssou D, Dubois D, Pirlot M, Prade H (eds) *Concepts and methods of decision-making process*. Wiley, New York, pp 771–812
- Inada K (1969) The simple majority decision rule. *Econometrica* 37:490–506
- Johnson DS (1990) A catalog of complexity classes. In: van Leeuwen J (ed) *Handbook of theoretical computer science*, vol A, Algorithms and complexity. Elsevier, Amsterdam, pp 67–161
- Johnson PE (1998) *Social choice theory and research*, CA, vol 123, Quantitative applications in the social sciences. Sage, Thousand Oaks
- Jünger M (1985) Polyhedral combinatorics and the acyclic subdigraph problem. Heldermann, Berlin
- Karp RM (1972) Reducibility among combinatorial problems. In: Miller RE, Thatcher JW (eds) *Complexity of computer computations*. Plenum Press, New York, pp 85–103
- Kelly JS (1987) *Social choice theory: an introduction*. Springer, Berlin
- Kemeny JG (1959) Mathematics without numbers. *Daedalus* 88:571–591
- Khachiyan L (1979) A polynomial algorithm in linear programming. *Sov Math Dokl* 20:191–194
- Köhler G (1978) *Choix multicritère et analyse algébrique de données ordinales*. Ph D thesis, université scientifique et médicale de Grenoble
- Lang J (2004) Logical preference representation and combinatorial vote. *Ann Math Artif Intell* 42:37–71
- Laslier J-F (1997) *Tournament solutions and majority voting*. Springer, Berlin
- Laslier J-F (2004) *Le vote et la règle majoritaire*. Analyse mathématique de la politique éditions du CNRS
- LeGrand R, Markakis E, Mehta A (2006) Approval voting: local search heuristics and approximation algorithms for the minimax solution. In: *Proceedings of the first international workshop on computational social choice (COMSOC 2006)*, University of Amsterdam, pp 234–247
- Levenglick A (1975) Fair and reasonable election systems. *Behav Sci* 20:34–46
- Levin J, Nalebuff B (1995) An introduction to vote-counting schemes. *J Econ Perspect* 9(1):3–26
- Lhuillier S (1794) Examen du mode d'élection proposé à la Convention nationale de France en février 1793 et adopté à Genève, Genève. Reprint In: (1976) *Math Sci Hum* 54:7–24
- Lines M (1986) Approval voting and strategy analysis: a venetian. *Ex Theor Decis* 20:155–172
- Mascart J (1919) *La vie et les travaux du chevalier Jean-Charles de Borda (1733–1799): épisodes de la vie scientifique au XVIIIe siècle*. Annales de l'université de Lyon vol. II (33). New edition, Presses de l'université de Paris-Sorbonne, 2000
- Maus S, Peters H, Storcken T (2006) Anonymous voting and minimal manipulability. In: *Proceedings of the first international workshop on computational social choice (COMSOC 2006)*, University of Amsterdam, pp 317–330
- Mc Garvey D (1953) A theorem on the construction of voting paradoxes. *Econometrica* 21:608–610
- McCabe-Dansted J (2006) Feasibility and approximability of Dodgson's rule. Master's thesis, University of Auckland
- McCabe-Dansted J, Pritchard G, Slinko A (2006) Approximability of Dodgson's rule. In: *Proceedings of the first international workshop on computational social choice (COMSOC 2006)*, University of Amsterdam, pp 234–247
- McKey B (2013) <http://cs.anu.edu.au/pp~bdm/data/digraphs.html>
- McLean I (1995) The first golden age of social choice, 1784–1803. In: Barnett WA, Moulin H, Salles M, Schofield NJ (eds) *Social choice, welfare, and ethics: proceedings of the eighth international symposium in economic theory and econometrics*. Cambridge University Press, Cambridge, pp 13–33
- McLean I, Hewitt F (1994) *Condorcet: foundations of social choice and political theory*. Edward Elgar, Hants
- McLean I, Urken A (1995) *Classics of social choice*. University of Michigan Press, Ann Arbor
- McLean I, Urken A (1997) *La réception des œuvres de Condorcet sur le choix social (1794–1803)*: Lhuillier, Morales et Daunou, in *Condorcet, Homme des Lumières et de la Révolution*, Chouillet A-M, Pierre Crépel (eds) ENS éditions, Fontenay-aux-roses, pp 147–160
- McLean I, McMillan A, Monroe BL (1995) Duncan Black and Lewis Carroll. *J Theor Polit* 7:107–124
- McLean I, Lorrey H, Colomer JM (2007) *Social choice in medieval Europe*. Workshop Histoire des Mathématiques Sociales, Paris
- Merrill S III, Grofman B (1999) *A unified theory of voting*. Cambridge University Press, Cambridge
- Miller N (1980) A new solution set for tournaments and majority voting: further graph-theoretical approaches to the theory of voting. *Am J Polit Sci* 24(1):68–96
- Mitlöhner J, Eckert D, Klamler C (2006) Simulating the effects of misperception on the manipulability of voting rules. In: *Proceedings of the first international workshop on computational social choice (COMSOC 2006)*, University of Amsterdam, p 234–247
- Monjardet B (1976) Lhuillier contre Condorcet au pays des paradoxes. *Math Sci Hum* 54:33–43
- Monjardet B (1979) Relations à éloignement minimum de relations binaires, note bibliographique. *Math Sci Hum* 67:115–122

- Monjardet B (1990) Sur diverses formes de la "règle de Condorcet" d'agrégation des préférences. *Math Inf Sci Hum* 111:61–71
- Monjardet B (2008a) Acyclic domains of linear orders: a survey. In: Brams S, Gehrlein WV, Roberts FS (eds) *The mathematics of preference, choice and order, essays in honor of Peter C. Fishburn*. Springer, Berlin, pp 139–160
- Monjardet B (2008b) *Mathématique Sociale and Mathematics*. A case study: Condorcet's effect and medians. *Electron J Hist Probab Stat* 4(1):1–26
- Moon JW (1968) *Topics on tournaments*. Holt, Rinehart and Winston, New York
- Morales JI (1797) *Memoria matemática sobre el cálculo de la opinión en las elecciones*. Imprenta Real, Madrid. Translated in McLean I, Urken A (1995) *Classics of social choice*. University of Michigan Press, Ann Arbor
- Moulin H (1980) On strategy-proofness and single peakedness. *Public Choice* 35:437–455
- Moulin H (1983) *The strategy of social choice*. North Holland, Amsterdam
- Moulin H (1985) Fairness and strategy in voting. In: Young HP (ed) *Fair allocation*, American Mathematical Society. *Proc Symp Appl Math* 33:109–142
- Moulin H (1986) Choosing from a tournament. *Soc Choice Welf* 3:272–291
- Nanson EJ (1882) Methods of election. *Trans Proc R Soc Vic* 18:197–240
- Nurmi H (1987) Comparing voting systems. D. Reidel, Dordrecht
- Pattanaik PK, Salles M (eds) (1983) *Social choice and welfare*. North-Holland, Amsterdam
- Poljak S, Turzik D (1986) A polynomial time heuristic for certain subgraph optimization problems with guaranteed lower bound. *Discret Math* 58:99–104
- Poljak S, Rödl V, Spencer J (1988) Tournament ranking with expected profit in polynomial time. *SIAM J Discret Math* 1(3):372–376
- Procaccia A, Rosenschein J (2006) Junta distribution and the average-case complexity of manipulating elections. In: *Proceedings of the 5th international joint autonomous agents and multiagent systems*, ACM Press, pp 497–504
- Procaccia A, Rosenschein J, Zohar A (2006) Multi-winner elections: complexity of manipulation, control, and winner-determination. In: *Proceedings of the 8th Trading Agent Design and Analysis and Agent Mediated Electronic Commerce Joint International workshop (TADA/AMEC 2006)*, pp 15–28
- Laplace (marquis de) PS (1795) *Journal de l'École Polytechnique*, tome II vol. 7–8; *Théorie analytique des probabilités*. Essai philosophique sur les probabilités. *Œuvres de Laplace*, tome VII, Paris, 1847
- Raman V, Saurabh S (2006) Parameterized algorithms for feedback set problems and their duals in tournaments. *Theor Comput Sci* 351:446–458
- Reid KB (2004) Tournaments. In: Gross JL, Yellen J (eds) *Handbook of graph theory*. CRC Press, Boca Raton, pp 156–184
- Reid KB, Beineke LW (1978) Tournaments. In: Beineke LW, Wilson RJ (eds) *Selected topics in graph theory*. Academic, London, pp 169–204
- Reinelt G (1985) *The linear ordering problem: algorithms and applications*, vol 8, Research and exposition in mathematics. Heldermann, Berlin
- Rothe J, Spakowski H (2006) On determining Dodgson winners by frequently self-knowingly correct algorithms and in average-case polynomial time. In: *Proceedings of the first international workshop on computational social choice (COMSOC 2006)*, University of Amsterdam, pp 234–247
- Rothe J, Spakowski H, Vogel J (2003) Exact complexity of the winner problem for Young elections. *Theor Comput Syst* 36(4):375–386
- Rowley CK (ed) (1993) *Social choice theory*, vol 1, The aggregation of preferences. Edward Elgar, London
- Saari D (1990) Susceptibility to manipulation. *Public Choice* 64:21–41
- Saari D (2001) *Decisions and elections, explaining the unexpected*. Cambridge University Press, Cambridge
- Satterthwaite M (1975) Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. *J Econ Theor* 10:187–217
- Schwartz T (1990) Cyclic tournaments and cooperative majority voting: a solution. *Soc Choice Welf* 7:19–29
- Simpson PB (1969) On defining areas of voter choice. *Q J Econ* 83(3):478–490
- Slater P (1961) Inconsistencies in a schedule of paired comparisons. *Biometrika* 48:303–312
- Smith JH (1973) Aggregation of preferences with variable electorate. *Econometrica* 41(6):1027–1041
- Smith D (1999) Manipulability measures of common social choice functions. *Soc Choice Welf* 16:639–661
- Spencer J (1971) Optimal ranking of tournaments. *Networks* 1:135–138
- Spencer J (1978) Nonconstructive methods in discrete mathematics. In: Rota GC (ed) *Studies in combinatorics*. Mathematical Association of America, Washington, DC, pp 142–178
- Spencer J (1987) Ten lectures on the probabilistic method. CBMS-NSF regional conference series in applied mathematics N 52, SIAM, Philadelphia
- Stearns R (1959) The voting problem. *Am Math Mon* 66:761–763
- Straffin PD Jr (1980) *Topics in the theory of voting*. Birkhäuser, Boston
- Taylor AD (1995) *Mathematics and politics strategy, voting, power, and proof*. Springer, Berlin
- Taylor AD (2005) *Social choice and the mathematics of manipulation*. Cambridge University Press, Cambridge
- Tideman TN (1987) Independence of clones as criterion for voting rules. *Soc Choice Welf* 4:185–206

- van Zuylen A (2005) Deterministic approximation algorithms for ranking and clusterings. Cornell ORIE technical report No. 1431
- Vazirani VV (2003) Approximation algorithms. Springer, Berlin
- Wakabayashi Y (1986) Aggregation of binary relations: algorithmic and polyhedral investigations. Ph D thesis, Augsburg
- Wakabayashi Y (1998) The complexity of computing medians of relations. *Resenhas* 3(3):323–349
- Weber RJ (1995) Approval voting. *J Econ Perspect* 9(1):39–49
- Woeginger GJ (2003) Banks winner in tournaments are difficult to recognize. *Soc Choice Welf* 20:523–528
- Young HP (1977) Extending Condorcet's rule. *J Econ Theor* 16(2):335–353
- ### Books and Reviews
- Aleskerov FT (1999) Arrovian aggregation models, mathematical and statistical methods, vol 39, Theory and decision library. Kluwer, Boston
- Aleskerov FT, Monjardet B (2002) Utility maximisation, choice and preference. Springer, Berlin
- Baker KM (1975) Condorcet from natural philosophy to social mathematics. The University of Chicago Press, Chicago. Reissued 1982
- Balinski M, Young HP (1982) Fair representation. Yale University Press, New Haven
- Barthélemy J-P, Monjardet B (1988) The median procedure in data analysis: new results and open problems. In: Bock HH (ed) Classification and related methods of data analysis. North Holland, Amsterdam
- Batteau P, Jacquet-Lagrèze É, Monjardet B (eds) (1981) Analyse et agrégation des préférences dans les sciences économiques et de gestion. Economica, Paris
- Black D (1996) Formal contributions to the theory of public choice. In: Brady GL, Tullock G (eds) The unpublished works of Duncan Black. Kluwer, Boston
- Bouyssou D, Marchant T, Pirlot M, Tsoukias A, Vincke P (2006) Evaluation and decision models with multiple criteria. Springer, Berlin
- Campbell DE (1992) Equity, efficiency, and social choice. Clarendon, Oxford
- Coughlin P (1992) Probabilistic voting theory. Cambridge University Press, Cambridge
- Danilov V, Sotskov A (2002) Social choice mechanisms. Springer, Berlin
- Dubois D, Pirlot M, Bouyssou D, Prade H (eds) (2006) Concepts et méthodes pour l'aide à la décision. Hermès, Paris
- Endriss U, Lang J (eds) (2006) Proceedings of the first international workshop on computational social choice, COMSOC 2006, University of Amsterdam
- Enelow J, Hinich M (eds) (1990) Advances in the spatial theory of voting. Cambridge University Press, Cambridge
- Farquharson R (1969) Theory of voting. Yale University Press, New Haven
- Feldman AM (1980) Welfare economics and social choice theory. Martinus Nijhoff, Boston
- Felsenthal DS, Machover M (1998) The measurement of voting power: theory and practice, problems and paradoxes. Edward Elgar, Cheltenham
- Gaertner W (2001) Domains conditions in social choice theory. Cambridge University Press, Cambridge
- Greenberg J (1990) The theory of social situations. Cambridge University Press, Cambridge
- Grofman B (1981) When is the Condorcet winner the Condorcet winner? University of California, Irvine
- Grofman B, Owen G (eds) (1986) Information pooling and group decision making. JAI Press, Greenwich
- Heal G (ed) (1997) Topological social choice. Springer, Berlin
- Hillinger C (2004) Voting and the cardinal aggregation cardinal of judgments. Discussion papers in economics 353, University of Munich
- Holler MJ (ed) (1978) Power voting and voting power. Physica, Wurtsburg
- Holler MJ, Owen G (eds) (2001) Indices and coalition formation. Kluwer, Boston
- Kemeny J, Snell L (1960) Mathematical models in the social sciences. Ginn, Boston
- Laslier J-F (2006) Spatial approval voting. *Polit Anal* 14(2):160–185
- Laslier J-F, Van Der Straeten K (2008) A live experiment on approval voting. *Exp Econ* 11:97–105
- Lieberman B (ed) (1971) Social choice. Gordon and Breach, New York
- Mirkin BG (1979) Group choice. Winston, Washington, DC
- Moulin H (2003) Fair division and collective welfare. Institute of Technology Press, Boston
- Nurmi H (1999) Voting paradoxes and how to deal with them. Springer, Berlin
- Nurmi H (2002) Voting procedures under uncertainty. Springer, Berlin
- Pattanaik PK (1971) Voting and collective choice. Harvard University Press, Cambridge
- Pattanaik PK (1978) Strategy and group choice. North Holland, Amsterdam
- Peleg B (1984) Game theoretic analysis of voting in committees. Cambridge University Press, Cambridge
- Peleg B, Peters H (2010) Strategic social choice. Springer, Berlin
- Rothschild E (2001) Economic sentiments: Adam Smith, Condorcet, and the enlightenment. Harvard University Press, Cambridge
- Saari DG (1994) Geometry of voting. Springer, Berlin
- Saari DG (1995) Basic geometry of voting. Springer, Berlin
- Saari DG (2000) Chaotic elections! American Mathematical Society, Providence
- Schofield N (1984) Social choice and democracy. Springer, Berlin

- Schofield N (ed) (1996) Collective decision making: social choice and political economy. Kluwer, Boston
- Schwartz T (1986) The logic of collective choice. Columbia University Press, New York
- Sen AK (1979) Collective choice and social welfare. North Holland, Amsterdam
- Sen AK (1982) Choice, welfare and measurement. Basil Blackwell, Oxford
- Suzumura K (1984) Rational choice, collective decisions and social welfare. Cambridge University Press, Cambridge
- Tanguiane AS (1991) Aggregation and representation of preferences, introduction to mathematical theory of democracy. Springer, Berlin
- Tideman N (2006) Collective decisions and voting: the potential for public choice. Ashgate, Burlington
- van Deemen A, Rusinowska A (eds) (2010) Collective decision making. Springer, Berlin
- Woodall DR (1997) Monotonicity of single-seat preferential election rules. *Discret Appl Math* 77:81–98
- Young HP (1974) An axiomatization of Borda's rule. *J Econ Theor* 9:43–52
- Young HP (1986) Optimal ranking and choice from pairwise comparisons. In: Grofman B, Owen G (eds) *Information pooling and group decision making*. JAI Press, Greenwich, pp 113–122
- Young HP (1988) Condorcet theory of voting. *Am Polit Sci Rev* 82:1231–1244
- Young HP (1995) Optimal voting rules. *J Econ Perspect* 9(1):51–64
- Young HP, Levenglick A (1978) A consistent extension of Condorcet's election principle. *SIAM J Appl Math* 35:285–300

Evolutionary Game Theory

William H. Sandholm

Department of Economics, University of Wisconsin, Madison, USA

Article Outline

Glossary

Definition of the Subject

Introduction

Normal Form Games

Static Notions of Evolutionary Stability

Population Games

Revision Protocols

Deterministic Dynamics

Stochastic Dynamics

Local Interaction

Applications

Future Directions

Bibliography

Glossary

Deterministic evolutionary dynamic A deterministic evolutionary dynamic is a rule for assigning population games to ordinary differential equations describing the evolution of behavior in the game. Deterministic evolutionary dynamics can be derived from *revision protocols*, which describe choices (in economic settings) or births and deaths (in biological settings) on an agent-by-agent basis.

Evolutionarily stable strategy (ESS) In a symmetric normal form game, an evolutionarily stable strategy is a (possibly mixed) strategy with the following property: a population in which all members play this strategy is resistant to invasion by a small group of mutants who play an alternative mixed strategy.

Normal form game A normal form game is a strategic interaction in which each of n players chooses a strategy and then receives a payoff that depends on all agents' choices of strategy. In a *symmetric two-player normal form game*, the two players choose from the same set of strategies, and payoffs only depend on own and opponent's choices, not on a player's identity.

Population game A population game is a strategic interaction among one or more large populations of agents. Each agent's payoff depends on his own choice of strategy and the distribution of others' choices of strategies. One can generate a population game from a normal form game by introducing random matching; however, many population games of interest, including congestion games, do not take this form.

Replicator dynamic The replicator dynamic is a fundamental deterministic evolutionary dynamic for games. Under this dynamic, the percentage growth rate of the mass of agents using each strategy is proportional to the excess of the strategy's payoff over the population's average payoff. The replicator dynamic can be interpreted biologically as a model of natural selection, and economically as a model of imitation.

Revision protocol A revision protocol describes both the timing and the results of agents' decisions about how to behave in a repeated strategic interaction. Revision protocols are used to derive both deterministic and stochastic evolutionary dynamics for games.

Stochastically stable state In Game-theoretic models of stochastic evolution in games are often described by irreducible Markov processes. In these models, a population state is stochastically stable if it retains positive weight in the process's stationary distribution as the level of noise in agents' choices approaches zero, or as the population size approaches infinity.

Definition of the Subject

Evolutionary game theory studies the behavior of large populations of agents who repeatedly engage in strategic interactions. Changes in behavior in these populations are driven either by natural selection via differences in birth and death rates, or by the application of myopic decision rules by individual agents.

The birth of evolutionary game theory is marked by the publication of a series of papers by mathematical biologist John Maynard Smith (1972, 1974; Maynard Smith and Price 1973). Maynard Smith adapted the methods of traditional game theory (Nash 1951; von Neumann and Morgenstern 1944), which were created to model the behavior of rational economic agents, to the context of biological natural selection. He proposed his notion of an *evolutionarily stable strategy (ESS)* as a way of explaining the existence of ritualized animal conflict.

Maynard Smith's equilibrium concept was provided with an explicit dynamic foundation through a differential equation model introduced by Taylor and Jonker (1978). Schuster and Sigmund (1983), following Dawkins (1976), dubbed this model the *replicator dynamic*, and recognized the close links between this game-theoretic dynamic and dynamics studied much earlier in population ecology (Lotka 1920; Volterra 1931) and population genetics (Fisher 1930). By the 1980s, evolutionary game theory was a well-developed and firmly established modeling framework in biology (Hofbauer and Sigmund 1988).

Towards the end of this period, economists realized the value of the evolutionary approach to game theory in social science contexts, both as a method of providing foundations for the equilibrium concepts of traditional game theory, and as a tool for selecting among equilibria in games that admit more than one. Especially in its early stages, work by economists in evolutionary game theory hewed closely to the interpretation set out by biologists, with the notion of ESS and the replicator dynamic understood as modeling natural selection in populations of agents genetically programmed to behave in specific ways. But it soon became clear that models of essentially the same form could be used to study the behavior of

populations of active decision makers (Crawford 1991; Friedman 1991; Mailath 1992; Nachbar 1990; Samuelson 1988; Selten 1991). Indeed, the two approaches sometimes lead to identical models: the replicator dynamic itself can be understood not only as a model of natural selection, but also as one of imitation of successful opponents (Björnerstedt and Weibull 1996; Schlag 1998; Weibull 1995).

While the majority of work in evolutionary game theory has been undertaken by biologists and economists, closely related models have been applied to questions in a variety of fields, including transportation science (Monderer and Shapley 1996; Nagurney and Zhang 1997; Sandholm 2001b, 2003, 2005b; Smith 1984), computer science (Fischer and Vöcking 2006; Sandholm 2001b, 2005b), and sociology (Bisin and Verdier 2001; Dokumaci and Sandholm 2007a; Kuran and Sandholm 2008; Zhang 2004a, b). Some paradigms from evolutionary game theory are close relatives of certain models from physics, and so have attracted the attention of workers in this field (Miękisz 2004; Szabó and Fáth 2007; Szabó and Hauert 2002; Tainaka 2001). All told, evolutionary game theory provides a common ground for workers from a wide range of disciplines.

Introduction

This article offers a broad survey of the theory of evolution in games. Section “[Normal Form Games](#)” introduces normal form games, a simple and commonly studied model of strategic interaction. Section “[Static Notions of Evolutionary Stability](#)” presents the notion of an evolutionarily stable strategy, a static definition of stability proposed for this normal form context.

Section “[Population Games](#)” defines population games, a general model of strategic interaction in large populations. Section “[Revision Protocols](#)” offers the notion of a revision protocol, an individual-level description of behavior used to define the population-level processes of central concern.

Most of the article concentrates on these population-level processes: section “[Deterministic Dynamics](#)” considers deterministic differential

equation models of game dynamics; section “[Stochastic Dynamics](#)” studies stochastic models of evolution based on Markov processes; and section “[Local Interaction](#)” presents deterministic and stochastic models of local interaction. Section “[Applications](#)” records a range of applications of evolutionary game theory, and section “[Future Directions](#)” suggests directions for future research. Finally, section “[Bibliography](#)” offers an extensive list of primary references.

Normal Form Games

In this section, we introduce a very simple model of strategic interaction: the symmetric two-player normal form game. We then define some of the standard solution concepts used to analyze this model, and provide some examples of games and their equilibria. With this background in place, we turn in subsequent sections to evolutionary analysis of behavior in games.

In a *symmetric two-player normal form game*, each of the two players chooses a (pure) strategy from the finite set S , which we write generically as $S = \{1, \dots, n\}$. The game’s *payoffs* are described by the matrix $A \in \mathbf{R}^n \times \mathbf{R}^n$. Entry A_{ij} is the payoff a player obtains when he chooses strategy i and his opponent chooses strategy j ; this payoff does not depend on whether the player in question is called player 1 or player 2.

The fundamental solution concept of noncooperative game theory is Nash equilibrium (Nash 1951). We say that the pure strategy $i \in S$ is a *symmetric Nash equilibrium* of A if

$$A_{ii} \geq A_{ji} \quad \text{for all } j \in S. \quad (1)$$

Thus, if his opponent chooses a symmetric Nash equilibrium strategy i , a player can do no better than to choose i himself.

A stronger requirement on strategy i demands that it be superior to all other strategies regardless of the opponent’s choice:

$$A_{ik} > A_{jk} \quad \text{for all } j, k \in S. \quad (2)$$

When condition (2) holds, we say that strategy i is *strictly dominant* in A .

Example 1 The game below, with strategies C (“cooperate”) and D (“defect”), is an instance of a *Prisoner’s Dilemma*:

	C	D
C	2 0	
D	3 1	

(To interpret this game, note that $A_{CD} = 0$ is the payoff to cooperating when one’s opponent defects.) Since $1 > 0$, defecting is a symmetric Nash equilibrium of this game. In fact, since $3 > 2$ and $1 > 0$, defecting is even a strictly dominant strategy. But since $2 > 1$, both players are better off when both cooperate than when both defect.

In many instances, it is natural to allow players to choose *mixed* (or *randomized*) *strategies*. When a player chooses mixed strategy from the simplex $X = \{x \in \mathbf{R}_+^n : \sum_{i \in S} x_i = 1\}$, his behavior is stochastic: he commits to playing pure strategy $i \in S$ with probability x_i .

When either player makes a randomized choice, we evaluate payoffs by taking expectations: a player choosing mixed strategy x against an opponent choosing mixed strategy y garners an expected payoff of

$$x' Ay = \sum_{i \in S} \sum_{j \in S} x_i A_{ij} y_j. \quad (3)$$

In biological contexts, payoffs are *fitnesses*, and represent levels of reproductive success relative to some baseline level; Eq. (3) reflects the idea that in a large population, expected reproductive success is what matters. In economic contexts, payoffs are *utilities*: a numerical representation of players’ preferences under which Eq. (3) captures players’ choices between uncertain outcomes (von Neumann and Morgenstern 1944).

The notion of Nash equilibrium extends easily to allow for mixed strategies. Mixed strategy x is a *symmetric Nash equilibrium* of A if

$$x' Ax \geq y' Ax \quad \text{for all } y \in X. \quad (4)$$

In words, x is a symmetric Nash equilibrium if its expected payoff against itself is at least as high

as the expected payoff obtainable by any other strategy y against x . Note that we can represent the pure strategy $i \in S$ using the mixed strategy $e_i \in X$, the i th standard basis vector in $\mathbf{R}^n$. If we do so, then definition (4) restricted to such strategies is equivalent to definition (1).

We illustrate these ideas with a few examples.

Example 2 Consider the *Stag Hunt* game:

$$\begin{array}{cc} & \begin{array}{cc} H & S \end{array} \\ \begin{array}{c} H \\ S \end{array} & \begin{array}{|cc|} \hline h & h \\ 0 & s \end{array} \end{array}.$$

Each player in the Stag Hunt game chooses between hunting hare (H) and hunting stag (S). A player who hunts hare always catches one, obtaining a payoff of $h > 0$. But hunting stag is only successful if both players do so, in which case each obtains a payoff of $s > h$. Hunting stag is potentially more profitable than hunting hare, but requires a coordinated effort.

In the Stag Hunt game, H and S (or, equivalently, e_H and e_S) are symmetric pure Nash equilibria. This game also has a symmetric mixed Nash equilibrium, namely $x^* = (x_H^*, x_S^*) = (\frac{s-h}{s}, \frac{h}{s})$. If a player's opponent chooses this mixed strategy, the player's expected payoff is h whether he chooses H, S, or any mixture between the two; in particular, x^* is a best response against itself.

To distinguish between the two pure equilibria, we might focus on the one that is *payoff dominant*, in that it achieves the higher joint payoff. Alternatively, we can concentrate on the *risk dominant* equilibrium (Harsanyi and Selten 1988), which utilizes the strategy preferred by a player who thinks his opponent is equally likely to choose either option (that is, against an opponent playing mixed strategy $(x_H, x_S) = (\frac{1}{2}, \frac{1}{2})$). In the present case, since $s > h$, equilibrium S is payoff dominant. Which strategy is risk dominant depends on further information about payoffs. If $s > 2h$, then S is risk dominant. But if $s < 2h$, H is risk dominant: evidently, payoff dominance and risk dominance need not agree.

Example 3 In the *Hawk-Dove* game (Maynard Smith 1982), the two players are animals

contesting a resource of value $v > 0$. The players choose between two strategies: display (D) or escalate (E). If both display, the resource is split; if one escalates and the other displays, the escalator claims the entire resource; if both escalate, then each player is equally likely to claim the entire resource or to be injured, suffering a cost of $c > v$ in the latter case.

The payoff matrix for the Hawk-Dove game is therefore

$$\begin{array}{cc} & \begin{array}{cc} D & E \end{array} \\ \begin{array}{c} D \\ E \end{array} & \begin{array}{|cc|} \hline \frac{1}{2}v & 0 \\ v & \frac{1}{2}(v-c) \end{array} \end{array}.$$

This game has no symmetric Nash equilibrium in pure strategies. It does, however, admit the symmetric mixed equilibrium $x^* = (x_D^*, x_E^*) = (\frac{c-v}{c}, \frac{v}{c})$. (In fact, it can be shown that every symmetric normal form game admits at least one symmetric mixed Nash equilibrium (Nash 1951).)

In this example, our focus on symmetric behavior may seem odd: rather than randomizing symmetrically, it seems more natural for players to follow an asymmetric Nash equilibrium in which one player escalates and the other displays. But the symmetric equilibrium is the most relevant one for understanding natural selection in populations whose members are randomly matched in pairwise contests -;see section “[Static Notions of Evolutionary Stability](#).”

Example 4 Consider the class of *Rock-Paper-Scissors* games:

$$\begin{array}{ccc} & \begin{array}{ccc} R & P & S \end{array} \\ \begin{array}{c} R \\ P \\ S \end{array} & \begin{array}{|ccc|} \hline 0 & -l & w \\ w & 0 & -l \\ -l & w & 0 \end{array} \end{array}.$$

Here $w > 0$ is the benefit of winning the match and $l > 0$ the cost of losing; ties are worth 0 to both players. We call this game *good RPS* if $w > l$, so that the benefit of winning the match exceeds the cost of losing, *standard RPS* if $w = l$, and *bad RPS* if $w < l$. Regardless of the values of w and l , the unique symmetric Nash equilibrium of this game, $x^* = (x_R^*, x_P^*, x_S^*) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, requires uniform randomization over the three strategies.

Static Notions of Evolutionary Stability

In introducing game-theoretic ideas to the study of animal behavior, Maynard Smith advanced this fundamental principle: that the evolutionary success of (the genes underlying) a given behavioral trait can depend on the prevalences of all traits. It follows that natural selection among the traits can be modeled as random matching of animals to play normal form games (Maynard Smith 1972, 1974, 1982; Maynard Smith and Price 1973). Working in this vein, Maynard Smith offered a stability concept for populations of animals sharing a common behavioral trait – that of playing a particular mixed strategy in the game at hand. Maynard Smith's concept of evolutionary stability, influenced by the work of Hamilton (1967) on the evolution of sex ratios, defines such a population as stable if it is resistant to invasion by a small group of mutants carrying a different trait.

Suppose that a large population of animals is randomly matched to play the symmetric normal form game A . We call mixed strategy $x \in X$ an *evolutionarily stable strategy (ESS)* if

$$x'A((1-\varepsilon)x + \varepsilon y) > y'A((1-\varepsilon)x + \varepsilon y) \quad (5)$$

for all $\varepsilon \leq \bar{\varepsilon}(y)$ and $y \neq x$.

To interpret condition (5), imagine that a population of animals programmed to play mixed strategy x is invaded by a group of mutants programmed to play the alternative mixed strategy y . Equation (5) requires that regardless of the choice of y , an incumbent's expected payoff from a random match in the post-entry population exceeds that of a mutant so long as the size of the invading group is sufficiently small.

The definition of ESS above can also be expressed as a combination of two conditions:

$$x'Ax \geq y'Ax \quad \text{for all } y \in X; \quad (4)$$

$$\begin{aligned} &\text{For all } y \neq x, \\ &[x'Ax = y'Ax] \text{ implies that } [x'Ay > y'Ay]. \end{aligned} \quad (6)$$

Condition (4) is familiar: it requires that the incumbent strategy x be a best response to itself, and so is none other than our definition of symmetric Nash equilibrium. Condition (6) requires

that if a mutant strategy y is an alternative best response against the incumbent strategy x , then the incumbent earns a higher payoff against the mutant than the mutant earns against itself.

A less demanding notion of stability can be obtained by allowing the incumbent and the mutant in condition (6) to perform equally well against the mutant:

$$\begin{aligned} &\text{For all } y \in X, \\ &[x'Ax = y'Ax] \text{ implies that } [x'Ay \geq y'Ay]. \end{aligned} \quad (7)$$

If x satisfies conditions (4) and (7), it is called a *neutrally stable strategy (NSS)* (Maynard Smith 1982).

Let us apply these stability notions to the games introduced in the previous section. Since every ESS and NSS must be a Nash equilibrium, we need only consider whether the Nash equilibria of these games satisfy the additional stability conditions, (6) and (7).

Example 5 In the Prisoner's Dilemma game (Example 1), the dominant strategy D is an ESS.

Example 6 In the Stag Hunt game (Example 2), each pure Nash equilibrium is an ESS. But the mixed equilibrium $(x_H^*, x_S^*) = (\frac{s-h}{s}, \frac{h}{s})$ is not an ESS: if mutants playing either pure strategy enter the population, they earn a higher payoff than the incumbents in the post-entry population.

Example 7 In the Hawk-Dove game (Example 3), the mixed equilibrium $(x_D^*, x_E^*) = (\frac{c-v}{c}, \frac{v}{c})$ is an ESS. Maynard Smith used this and other examples to explain the existence of ritualized fighting in animals. While an animal who escalates always obtains the resource when matched with an animal who merely displays, a population of escalators is unstable: it can be invaded by a group of mutants who display, or who merely escalate less often.

Example 8 In Rock-Paper-Scissors games (Example 4), whether the mixed equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is evolutionarily stable depends on the relative payoffs to winning and losing a match. In good RPS ($w > l$), x^* is an ESS; in standard RPS

($w = l$), x^* is a NSS but not an ESS, while in bad RPS ($w < l$), x^* is neither an ESS nor an NSS. The last case shows that neither evolutionary nor neutrally stable strategies need exist in a given game.

The definition of an evolutionarily stable strategy has been extended to cover a wide range of strategic settings, and has been generalized in a variety of directions. Prominent among these developments are set-valued versions of ESS: in rough terms, these concepts consider a set of mixed strategies $Y \subset X$ to be stable if the no population playing a strategy in the set can be invaded successfully by a population of mutants playing a strategy outside the set. Hines (1987) provides a thorough survey of the first 15 years of research on ESS and related notions of stability; key references on set-valued evolutionary solution concepts include (Balkenborg and Schlag 2001; Swinkels 1992; Thomas 1985).

Maynard Smith's notion of ESS attempts to capture the dynamic process of natural selection using a static definition. The advantage of this approach is that his definition is often easy to check in applications. Still, more convincing models of natural selection should be explicitly dynamic models, building on techniques from the theories of dynamical systems and stochastic processes. Indeed, this thoroughgoing approach can help us understand whether and when the ESS concept captures the notion of robustness to invasion in a satisfactory way.

The remainder of this article concerns explicitly dynamic models of behavior. In addition to being dynamic rather than static, these models will differ from the one considered in this section in two other important ways as well. First, rather than looking at populations whose members all play a particular mixed strategy, the dynamic models consider populations in which different members play different pure strategies. Second, instead of maintaining a purely biological point of view, our dynamic models will be equally well-suited to studying behavior in animal and human populations.

Population Games

Population games provide a simple and general framework for studying strategic interactions in large populations whose members play pure

strategies. The simplest population games are generated by random matching in normal form games, but the population game framework allows for interactions of a more intricate nature.

We focus here on games played by a single population (i.e., games in which all agents play equivalent roles). We suppose that there is a unit mass of agents, each of whom chooses a pure strategy from the set $S = \{1, \dots, n\}$. The aggregate behavior of these agents is described by a *population state* $x \in X$, with x_j representing the proportion of agents choosing pure strategy j . We identify a *population game* with a continuous vector-valued payoff function $F : X \rightarrow \mathbf{R}^n$. The scalar $F_i(x)$ represents the payoff to strategy i when the population state is x .

Population state x^* is a *Nash equilibrium* of F if no agent can improve his payoff by unilaterally switching strategies. More explicitly, x^* is a Nash equilibrium if

$$\begin{aligned} x_i^* > 0 & \text{ implies that } F_i(x) \\ & \geq F_j(x) \quad \text{for all } j \in S. \end{aligned} \quad (8)$$

Example 9 Suppose that the unit mass of agents are randomly matched to play the symmetric normal form game A . At population state x , the (expected) payoff to strategy i is the linear function $F_i(x) = \sum_{j \in S} A_{ij}x_j$; the payoffs to all strategies can be expressed concisely as $F(x) = Ax$. It is easy to verify that x^* is a Nash equilibrium of the population game F if and only if x^* is a symmetric Nash equilibrium of the symmetric normal form game A .

While population games generated by random matching are especially simple, many games that arise in applications are not of this form. In the biology literature, games outside the random matching paradigm are known as *playing the field* models (Maynard Smith 1982).

Example 10 Consider the following model of highway congestion (Beckmann et al. 1956; Monderer and Shapley 1996; Rosenthal 1973; Sandholm 2001b). A pair of towns, Home and Work, are connected by a network of *links*. To commute from Home to Work, an agent must choose a *path* $i \in S$ connecting the two towns. The payoff the agent obtains is the negation of the delay on the path he takes. The delay on the path is

the sum of the delays on its constituent links, while the delay on a link is a function of the number of agents who use that link.

Population games embodying this description are known as a *congestion games*. To define a congestion game, let Φ be the collection of links in the highway network. Each strategy $i \in S$ is a route from Home to Work, and so is identified with a set of links $\Phi_i \subseteq \Phi$. Each link ϕ is assigned a *cost function* $c_\phi : \mathbf{R}_+ \rightarrow \mathbf{R}$, whose argument is link ϕ 's *utilization level* u_ϕ :

$$u_\phi(x) = \sum_{i \in \rho(\phi)} x_i, \quad \text{where } \rho(\phi) = \{i \in S : \phi \in \Phi_i\}$$

The payoff of choosing route i is the negation of the total delays on the links in this route:

$$F_i(x) = - \sum_{\phi \in \Phi_i} c_\phi(u_\phi(x)).$$

Since driving on a link increases the delays experienced by other drivers on that link (i.e., since highway congestion involves *negative externalities*), cost functions in models of highway congestion are increasing; they are typically convex as well. Congestion games can also be used to model positive externalities, like the choice between different technological standards; in this case, the cost functions are decreasing in the utilization levels.

Revision Protocols

We now introduce foundations for our models of evolutionary dynamics. These foundations are built on the notion of a revision protocol, which describes both the timing and results of agents' myopic decisions about how to continue playing the game at hand (Benaïm and Weibull 2003; Björnerstedt and Weibull 1996; Hofbauer 1995a; Sandholm 2003; Weibull 1996). Revision protocols will be used to derive both the deterministic dynamics studied in section “[Deterministic Dynamics](#)” and the stochastic dynamics studied in section “[Stochastic Dynamics](#)”; similar ideas underlie the local interaction models introduced in section “[Local Interaction](#).”

Definition

Formally, a *revision protocol* is a map $\rho : \mathbf{R}^n \times X \rightarrow \mathbf{R}_+^{n \times n}$ that takes the payoff vectors π and population states x as arguments, and returns nonnegative matrices as outputs. For reasons to be made clear below, scalar $\rho_{ij}(\pi, x)$ is called the *conditional switch rate* from strategy i to strategy j .

To move from this notion to an explicit model of evolution, let us consider a population consisting of $N < \infty$ members. (A number of the analyzes to follow will consider the limit of the present model as the population size N approaches infinity – see sections “[Mean Dynamics](#),” “[Deterministic Approximation](#),” and “[Stochastic Stability via Large Population Limits](#).”) In this case, the set of feasible social states is the finite set $\mathcal{X}^N = X \cap \frac{1}{N} \mathbf{Z}^n = \{x \in X : Nx \in \mathbf{Z}^n\}$, a grid embedded in the simplex X .

A revision protocol ρ , a population game F , and a population size N define a continuous-time evolutionary process – a Markov process $\{X_t^N\}$ – on the finite state space $\mathcal{X}^N$. A one-size-fits-all description of this process is as follows. Each agent in the society is equipped with a “stochastic alarm clock.” The times between rings of an agent's clock are independent, each with a rate R exponential distribution. The ringing of a clock signals the arrival of a revision opportunity for the clock's owner. If an agent playing strategy $i \in S$ receives a revision opportunity, he switches to strategy $j \neq i$ with probability ρ_{ij}/R . If a switch occurs, the population state changes accordingly, from the old state x to a new state y that accounts for the agent's change in strategy.

While this interpretation of the evolutionary process can be applied to any revision protocol, simpler interpretations are sometimes available for protocols with additional structure. The examples to follow illustrate this point.

Examples

Imitation Protocols and Natural Selection Protocols

In economic contexts, revision protocols of the form

$$\rho_{ij}(\pi, x) = x_j \hat{\rho}_{ij}(\pi, x) \quad (9)$$

are called *imitation protocols* (Björnerstedt and Weibull 1996; Hofbauer 1995a; Weibull 1995).

These protocols can be given a very simple interpretation: when an agent receives a revision opportunity, he chooses an opponent at random and observes her strategy. If our agent is playing strategy i and the opponent strategy j , the agent switches from i to j with probability proportional to $\hat{\rho}_{ij}$. Notice that the value of the population share x_j is not something the agent need know; this term in (9) accounts for the agent's observing a randomly chosen opponent.

Example 11 Suppose that after selecting an opponent, the agent imitates the opponent only if the opponent's payoff is higher than his own, doing so in this case with probability proportional to the payoff difference:

$$\rho_{ij}(\pi, x) = x_j [\pi_j - \pi_i]_+.$$

This protocol is known as *pairwise proportional imitation* (Schlag 1998).

Protocols of form (9) also appear in biological contexts, (Moran 1962; Nowak 2006; Nowak et al. 2004), where in these cases we refer to them as *natural selection protocols*. The biological interpretation of (9) supposes that each agent is programmed to play a single pure strategy. An agent who receives a revision opportunity dies, and is replaced through asexual reproduction. The reproducing agent is a strategy j player with probability $\rho_{ij}(\pi, x) = x_j \hat{\rho}_{ij}(\pi, x)$, which is proportional both to the number of strategy j players and to some function of the prevalences and fitnesses of all strategies. Note that this interpretation requires the restriction

$$\sum_{j \in S} \rho_{ij}(\pi, x) \equiv 1.$$

Example 12 Suppose that payoffs are always positive, and let

$$\rho_{ij}(\pi, x) = \frac{x_j \pi_j}{\sum_{k \in S} x_k \pi_k}. \quad (10)$$

Understood as a natural selection protocol, (10) says that the probability that the reproducing agent is a strategy j player is proportional to $x_j \pi_j$, the aggregate fitness of strategy j players.

In economic contexts, we can interpret (10) as an imitative protocol based on repeated sampling.

When an agent's clock rings he chooses an opponent at random. If the opponent is playing strategy j , the agent imitates him with probability proportional to π_j . If the agent does not imitate this opponent, he draws a new opponent at random and repeats the procedure.

Direct Evaluation Protocols

In the previous examples, only strategies currently in use have any chance of being chosen by a revising agent (or of being the programmed strategy of the newborn agent). Under other protocols, agents' choices are not mediated through the population's current behavior, except indirectly via the effect of behavior on payoffs. These *direct evaluation protocols* require agents to directly evaluate the payoffs of the strategies they consider, rather than to indirectly evaluate them as under an imitative procedure.

Example 13 Suppose that choices are made according to the *logit choice rule*:

$$\rho_{ij}(\pi, x) = \frac{\exp(\eta^{-1} \pi_j)}{\sum_{k \in S} \exp(\eta^{-1} \pi_k)}. \quad (11)$$

The interpretation of this protocol is simple. Revision opportunities arrive at unit rate. When an opportunity is received by an i player, he switches to strategy j with probability $\rho_{ij}(\pi, x)$, which is proportional to an exponential function of strategy j 's payoffs. The parameter $\eta > 0$ is called the *noise level*. If η is large, choice probabilities under the logit rule are nearly uniform. But if η is near zero, choices are optimal with probability close to one, at least when the difference between the best and second best payoff is not too small.

Additional examples of revision protocols can be found in the next section, and one can construct new revision protocols by taking linear combinations of old ones; see (Sandholm 2017) for further discussion.

Deterministic Dynamics

Although antecedents of this approach date back to the early work of Brown and von Neumann (1950), the use of differential equations to model evolution in games took root with the introduction of the

replicator dynamic by Taylor and Jonker (1978), and remains an vibrant area of research; Hofbauer and Sigmund (2003) and Sandholm (2017) offer recent surveys. In this section, we derive a deterministic model of evolution: the *mean dynamic* generated by a revision protocol and a population game. We study this deterministic model from various angles, focusing in particular on local stability of rest points, global convergence to equilibrium, and nonconvergent limit behavior.

While the bulk of the literature on deterministic evolutionary dynamics is consistent with the approach we take here, we should mention that other specifications exist, including discrete time dynamics (Akin and Losert 1984; Dekel and Scotchmer 1992; Losert and Akin 1983; Weissing 1991), and dynamics for games with continuous strategy sets (Bomze 1990, 1991; Friedman and Yellin 1997; Hofbauer et al. 2005; Oechssler and Riedel 2001, 2002) and for Bayesian population games (Dokumaci and Sandholm 2007a; Ely and Sandholm 2005; Sandholm 2007a). Also, deterministic dynamics for extensive form games introduce new conceptual issues; see (Binmore et al. 1995a; Binmore and Samuelson 1999; Cressman 1996, 2000; Cressman and Schlag 1998) and the monograph of Cressman (2003).

Mean Dynamics

As described earlier in section “Definition,” a revision protocol ρ , a population game F , and a population size N define a Markov process $\{X_t^N\}$ on the finite state space $\mathcal{X}^N$. We now derive a deterministic process – the *mean dynamic* – that describes the expected motion of $\{X_t^N\}$. In section “Deterministic Approximation,” we will describe formally the sense in which this deterministic process provides a very good approximation of the behavior of the stochastic process $\{X_t^N\}$, at least over finite time horizons and for large population sizes. But having noted this result, we will focus in this section on the deterministic process itself.

To compute the expected increment of $\{X_t^N\}$ over the next dt time units, recall first that each of the N agents receives revision opportunities via a rate R exponential distribution, and so expects to receive Rdt opportunities during the next dt time units. If the current state is x , the expected number

of revision opportunities received by agents currently playing strategy i is approximately $Nx_i Rdt$. Since an i player who receives a revision opportunity switches to strategy j with probability ρ_{ij}/R , the expected number of such switches during the next dt time units is approximately $Nx_i \rho_{ij}dt$. Therefore, the expected change in the number of agents choosing strategy i during the next dt time units is approximately

$$N \left(\sum_{j \in S} x_j \rho_{ji}(F(x), x) - x_i \sum_{j \in S} \rho_{ij}(F(x), x) \right) dt. \quad (12)$$

Dividing expression (12) by N and eliminating the time differential dt yields a differential equation for the rate of change in the *proportion* of agents choosing strategy i :

$$\dot{x}_i = \sum_{j \in S} x_j \rho_{ji}(F(x), x) - x_i \sum_{j \in S} \rho_{ij}(F(x), x). \quad (M)$$

Equation (M) is the *mean dynamic* (or *mean field*) generated by revision protocol ρ in population game F . The first term in (M) captures the inflow of agents to strategy i from other strategies, while the second captures the outflow of agents to other strategies from strategy i .

Examples

We now describe some examples of mean dynamics, starting with ones generated by the revision protocols from section “Examples.” To do so, we let

$$\bar{F}(x) = \sum_{i \in S} x_i F_i(x)$$

denote the *average payoff* obtained by the members of the population, and define the *excess payoff* to strategy i ,

$$\hat{F}_i(x) = F_i(x) - \bar{F}(x),$$

to be the difference between strategy i ’s payoff and the population’s average payoff.

Example 14 In Example 11, we introduced the pairwise proportional imitation protocol $\rho_{ij}(\pi, x)$

$= x_j[\pi_j - \pi_i]_+$. This protocol generates the mean dynamic

$$\dot{x}_i = x_i \hat{F}_i(x). \quad (13)$$

Equation (13) is the *replicator dynamic* (Taylor and Jonker 1978), the best-known dynamic in evolutionary game theory. Under this dynamic, the percentage growth rate $\dot{x}_i/x_i$, of each strategy currently in use is equal to that strategy's current excess payoff; unused strategies always remain so. There are a variety of revision protocols other than pairwise proportional imitation that generate the replicator dynamic as their mean dynamics; see (Björnerstedt and Weibull 1996; Hofbauer 1995a; Hofbauer and Sigmund 2003; Weibull 1996).

Example 15 In Example 12, we assumed that payoffs are always positive, and introduced the protocol $\rho_{ij} \propto x_j \pi_j$, which we interpreted both as a model of biological natural selection and as a model of imitation with repeated sampling. The resulting mean dynamic,

$$\dot{x}_i = \frac{x_i F_i(x)}{\sum_{k \in S} x_k F_k(x)} - x_i = \frac{x_i \hat{F}_i(x)}{\bar{F}(x)}, \quad (14)$$

is the *Maynard Smith replicator dynamic* (Maynard Smith 1982). This dynamic only differs from the standard replicator dynamic (13) by a change of speed, with motion under (14) being relatively fast when average payoffs are relatively low. (In multipopulation models, the two dynamics are less similar, and convergence under one does not imply convergence under the other – see (Sandholm 2017; Weibull 1995).)

Example 16 In Example 13 we introduced the logit choice rule $\rho_{ij}(\pi, x) \propto \exp(\eta^{-1} \pi_j)$. The corresponding mean dynamic,

$$\dot{x}_j = \frac{\exp(\eta^{-1} F_j(x))}{\sum_{k \in S} \exp(\eta^{-1} F_k(x))} - x_j, \quad (15)$$

is called the *logit dynamic* (Fudenberg and Levine 1998).

If we take the noise level η to zero, then the probability with which a revising agent chooses the best response approaches one whenever the best response is unique. At such points, the logit dynamic approaches the *best response dynamic* (Gilboa and Matsui 1991):

$$\dot{x} \in B^F(x) - x, \quad (16)$$

where

$$B^F(x) = \operatorname{argmax}_{y \in X} y' F(x)$$

defines the (*mixed*) *best response correspondence* for game F . Note that unlike the other dynamics we consider here, (16) is defined not by an ordinary differential equation, but by a differential inclusion, a formulation proposed in Hofbauer (1995b).

Example 17 Consider the protocol

$$\rho_{ij}(\pi, x) = \left[\pi_j - \sum_{k \in S} x_k \pi_k \right]_+.$$

When an agent's clock rings, he chooses a strategy at random; if that strategy's payoff is above average, the agent switches to it with probability proportional to its excess payoff. The resulting mean dynamic,

$$\dot{x}_i BM = [\hat{F}_i(x)]_+ - x_i \sum_{k \in S} [\hat{F}_k(x)]_+,$$

is called the *Brown-von Neumann-Nash (BNN) dynamic* (Brown and von Neumann 1950); see also (Hofbauer 2000; Sandholm 2005a; Skyrms 1990; Swinkels 1993; Weibull 1996).

Example 18 Consider the revision protocol

$$\rho_{ij}(\pi, x) = [\pi_j - \pi_i]_+.$$

When an agent's clock rings, he selects a strategy at random. If the new strategy's payoff is higher than his current strategy's payoff, he switches strategies with probability proportional to the difference between the two payoffs. The resulting mean dynamic,

$$\dot{x}_i = \sum_{j \in S} x_j [F_i(x) - F_j(x)]_+ - x_i \sum_{j \in S} [F_j(x) - F_i(x)]_+, \quad (17)$$

is called the *Smith dynamic* (Smith 1984); see also (Sandholm 2006).

We summarize these examples of revision protocols and mean dynamics in Table 1.

Figure 1 presents phase diagrams for the five basic dynamics when the population is randomly matched to play standard Rock-Paper-Scissors (Example 4). In the phase diagrams, colors represent speed of motion: within each diagram, motion is fastest in the red regions and slowest in the blue ones.

The phase diagram of the replicator dynamic reveals closed orbits around the unique Nash equilibrium $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Since this dynamic is based on imitation (or on reproduction), each face and each vertex of the simplex X is an invariant set: a strategy initially absent from the population will never subsequently appear.

The other four dynamics pictured are based on direct evaluation, allowing agents to select strategies that are currently unused. In these cases, the Nash equilibrium is the sole rest point, and attracts solutions from all initial conditions. (In the case of the logit dynamic, the rest point happens to coincide with the Nash equilibrium only because of the symmetry of the game; see (Hofbauer and Sandholm 2002, 2007).) Under the logit and best response dynamics, solution trajectories quickly change direction and then accelerate when the

best response to the population state changes; under the BNN and especially the Smith dynamic, solutions approach the Nash equilibrium in a less angular fashion.

Evolutionary Justification of Nash Equilibrium

One of the goals of evolutionary game theory is to justify the prediction of Nash equilibrium play. For this justification to be convincing, it must be based on a model that makes only mild assumptions about agents' knowledge about one another's behavior. This sentiment can be captured by introducing two desiderata for revision protocols:

- (C)Continuity : ρ is Lipschitz continuous.
- (SD)Scarcity of data : ρ_{ij} only depends on π_i, π_j , and x_j .


Continuity (C) asks that revision protocols depend continuously on their inputs, so that small changes in aggregate behavior do not lead to large changes in players' responses. *Scarcity of data* (SD) demands that the conditional switch rate from strategy i to strategy j only depend on the payoffs of these two strategies, so that agents need only know those facts that are most germane to the decision at hand (Sandholm 2017). (The dependence of ρ_{ij} on x_j is included to allow for dynamics based on imitation.) Protocols that respect these two properties do not make unrealistic demands on the amount of information that agents in an evolutionary model possess.

Table 1 Five basic deterministic dynamics

Revision protocol	Mean dynamic	Name and source
$\rho_{ij} = x_j[\pi_j - \pi_i]_+$	$\dot{x}_i = x_i \bar{F}_i(x)$	Replicator (Taylor and Jonker 1978)
$\rho_{ij} = \frac{\exp(\eta^{-1}\pi_j)}{\sum_{k \in S} \exp(\eta^{-1}\pi_k)}$	$\dot{x}_i = \frac{\exp(\eta^{-1}F_i(x))}{\sum_{k \in S} \exp(\eta^{-1}F_k(x))} - x_i$	Logit (Fudenberg and Levine 1998)
$\rho_{ij} = 1_{\{j=\arg\max_{k \in S} \pi_k\}}$	$\dot{x} \in B^F(x) - x$	Best response (Gilboa and Matsui 1991)
$\rho_{ij} = [\pi_j - \sum_{k \in S} x_k \pi_k]_+$	$\dot{x}_i = [\bar{F}_i(x)]_+ - x_i \sum_{j \in S} [\bar{F}_j(x)]_+$	BNN (Brown and von Neumann 1950)
$\rho_{ij} = [\pi_j - \pi_i]_+$	$\dot{x}_i = \sum_{j \in S} x_j [F_i(x) - F_j(x)]_+$	Smith (1984)
	$-\dot{x}_i = \sum_{j \in S} x_j [F_j(x) - F_i(x)]_+$	

Evolutionary Game

Theory, Fig. 1 Five basic deterministic dynamics in standard Rock-Paper-Scissors. Colors represent speeds: *red* is fastest, *blue* is slowest



Our two remaining desiderata impose restrictions on mean dynamics $\dot{x} = V^F(x)$, linking the evolution of aggregate behavior to incentives in the underlying game.

(NS) Nash stationarity :

$$V^F(x) = \mathbf{0} \quad \text{if and only if} \quad x \in NE(F).(\text{PC})$$

Positive correlation :

$$V^F(x) \neq \mathbf{0} \quad \text{implies that} \quad V^F(x)'F(x) > 0.$$

Nash stationarity (NS) is a restriction on stationary states: it asks that the rest points of

the mean dynamic be precisely the Nash equilibria of the game being played. *Positive correlation* (PC) is a restriction on disequilibrium adjustment: it requires that away from rest points, strategies' growth rates be positively correlated with their payoffs. Condition (PC) is among the weakest of the many conditions linking growth rates of evolutionary dynamics and payoffs in the underlying game; for alternatives, see (Friedman 1991; Hofbauer and Weibull 1996; Nachbar 1990; Ritzberger

Evolutionary Game Theory, Table 2 Families of deterministic evolutionary dynamics and their properties; yes* indicates that a weaker or alternate form of the property is satisfied

Dynamic	Family	(C)	(SD)	(NS)	(PC)
Replicator	Imitation	yes	yes	no	yes
Best response		no	yes*	yes*	yes*
Logit	Perturbed best response	yes	yes*	no	no
BNN	Excess payoff	yes	no	yes	yes
Smith	Pairwise comparison	yes	yes	yes	yes

and Weibull 1995; Samuelson and Zhang 1992; Sandholm 2001b; Swinkels 1993).

In Table 2, we report how the five basic dynamics fare under the four criteria above. For the purposes of justifying the Nash prediction, the most important row in the table is the last one, which reveals that the Smith dynamic satisfies all four desiderata at once: while the revision protocol for the Smith dynamic (see Example 18) requires only limited information on the part of the agents who employ it, this information is enough to ensure that rest points of the dynamic and Nash equilibria coincide.

In fact, the dynamics introduced above can be viewed as members of families of dynamics that are based on similar revision protocols and that have similar qualitative properties. For instance, the Smith dynamic is a member of the family of *pairwise comparison* dynamics (Sandholm 2006), under which agents only switch to strategies that outperform their current choice. For this reason, the exact functional forms of the previous examples are not essential to establishing the properties noted above.

In interpreting these results, it is important to remember that Nash stationarity only concerns the rest points of a dynamic; it says nothing about whether a dynamic will converge to Nash equilibrium from an arbitrary initial state. The question of convergence is addressed in sections “[Global Convergence](#)” and “[Nonconvergence](#).” There we will see that in some classes of games, general guarantees of convergence can be obtained, but that there are some games in which no reasonable dynamic converges to equilibrium.

Local Stability

Before turning to the global behavior of evolutionary dynamics, we address the question of local

stability. As we noted at the onset, an original motivation for introducing game dynamics was to provide an explicitly dynamic foundation for Maynard Smith’s notion of ESS (Taylor and Jonker 1978). Some of the earliest papers on evolutionary game dynamics (Hofbauer et al. 1979; Zeeman 1980) established that being an ESS is a sufficient condition for asymptotically stability under the replicator dynamic, but that it is not a necessary condition. It is curious that this connection obtains despite the fact that ESS is a stability condition for a population whose members all play the same mixed strategy, while (the usual version of) the replicator dynamic looks at populations of agents choosing among different pure strategies.

In fact, the implications of ESS for local stability are not limited to the replicator dynamic. Suppose that the symmetric normal form game A admits a symmetric Nash equilibrium that places positive probability on each strategy in S . One can show that this equilibrium is an ESS if and only if the payoff matrix A is negative definite with respect to the tangent space of the simplex:

$$\begin{aligned} z'Az &< 0 \quad \text{for all } z \in TX \\ &= \left\{ \hat{z} \in \mathbf{R}^n : \sum_{i \in S} \hat{z}_i = 0 \right\}. \end{aligned} \quad (18)$$

Condition (18) and its generalizations imply local stability of equilibrium not only under the replicator dynamic, but also under a wide range of other evolutionary dynamics: see (Cressman 1997; Hofbauer 2000; Hofbauer and Hopkins 2005; Hofbauer and Sandholm 2006a; Hopkins 1999; Sandholm 2007a) for further details.

The papers cited above use linearization and Lyapunov function arguments to establish local stability. An alternative approach to local stability analysis, via index theory, allows one to establish

restrictions on the stability properties of all rest points at once – see (Demichelis and Ritzberger 2003).

Global Convergence

While analyses of local stability reveal whether a population will return to equilibrium after a small disturbance, they do not tell us whether the population will approach equilibrium from an arbitrary disequilibrium state. To establish such global convergence results, we must restrict attention to classes of games defined by certain interesting payoff structures. These structures appear in applications, lending strong support for the Nash prediction in the settings where they arise.

Potential Games

A *potential game* (Beckmann et al. 1956; Hofbauer and Sigmund 1988; Monderer and Shapley 1996; Rosenthal 1973; Sandholm 2001b, 2007c) is a game that admits a *potential function*: a scalar valued function whose gradient describes the game's payoffs. In a full potential game $F : \mathbf{R}_+^n \rightarrow \mathbf{R}^n$ (see Sandholm 2007c), all information about incentives is captured by the potential function $f : \mathbf{R}_+^n \rightarrow \mathbf{R}$, in the sense that

$$\nabla f(x) = F(x) \quad \text{for all } x \in \mathbf{R}_+^n. \quad (19)$$

If F is smooth, then it is a full potential game if and only if it satisfies *full externality symmetry*.

$$\frac{\partial F_i}{\partial x_j}(x) = \frac{\partial F_j}{\partial x_i}(x) \quad \text{for all } i, j \in S \text{ and } x \in \mathbf{R}_+^n. \quad (20)$$

That is, the effect on the payoff to strategy i of adding new strategy j players always equals the effect on the payoff to strategy j of adding new strategy i players.

Example 19 Suppose a single population is randomly matched to play the symmetric normal form game $A \in \mathbf{R}^{n \times n}$, generating the population game $F(x) = Ax$. We say that A exhibits *common interests* if the two players in a match always receive the same payoff. This means that $A_{ij} = A_{ji}$

for all i and j , or, equivalently, that the matrix A is symmetric. Since $DF(x) = A$, this is precisely what we need for F to be a full potential game. The full potential function for F is $f(x) = \frac{1}{2}x'Ax$, which is one-half of the average payoff function $\bar{F}(x) = \sum_{i \in S} x_i F_i(x) = x'Ax$. The common interest assumption defines a fundamental model from population genetics, this assumption reflects the shared fate of two genes that inhabit the same organism (Fisher 1930; Hofbauer and Sigmund 1988, 1998).

Example 20 In Example 10, we introduced congestion games, a basic model of network congestion. To see that these games are potential games, observe that an agent taking path $j \in S$ affects the payoffs of agents choosing path $i \in S$ through the marginal increases in congestion on the links $\phi \in \Phi_i \cap \Phi_j$ that the two paths have in common. But since the marginal effect of an agent taking path i on the payoffs of agents choosing path j is identical, full externality symmetry (20) holds:

$$\frac{\partial F_i}{\partial x_j}(x) = - \sum_{\phi \in \Phi_i \cap \Phi_j} c'_\phi(u_\phi(x)) = \frac{\partial F_j}{\partial x_i}(x).$$

In congestion games, the potential function takes the form

$$f(x) = - \sum_{\phi \in \Phi} \int_0^{u_\phi(x)} c_\phi(z) \, dz,$$

and so is typically unrelated to aggregate payoffs,

$$\bar{F}(x) = \sum_{i \in S} x_i F_i(x) = - \sum_{\phi \in \Phi} u_\phi(x) c_\phi(u_\phi(x)).$$

However, potential is proportional to aggregate payoffs if the cost functions c_ϕ are all monomials of the same degree (Dafermos and Sparrow 1969; Sandholm 2001b).

Population state x is a Nash equilibrium of the potential game F if and only if it satisfies the Kuhn-Tucker first order conditions for maximizing the potential function f on the simplex X (Beckmann et al. 1956; Sandholm 2001b). Furthermore, it is simple to verify that any dynamic

$\dot{x} = V^F(x)$ satisfying positive correlation (PC) ascends the potential function:

$$\frac{d}{dt}f(x_t) = \nabla f(x_t)'x_t = F(x_t)'V^F(x_t) \geq 0.$$


It then follows from classical results on Lyapunov functions that any dynamic satisfying positive correlation (PC) converges to a connected set of rest points. If the dynamic also satisfies Nash stationarity (NS), these sets consist entirely of Nash equilibria. Thus, in potential games, very mild conditions on agents' adjustment rules are sufficient to justify the prediction of Nash equilibrium play.

In the case of the replicator dynamic, one can say more. On the interior of the simplex X , the replicator dynamic for the potential game F is a *gradient system* for the potential function f (i.e., it always ascends f in the direction of maximum increase). However, this is only true after one introduces an appropriate Riemannian metric on X (Kimura 1958; Shahshahani 1979). An equivalent statement of this result, due to Akin (1979), is that the replicator dynamic is the gradient system for f under the usual Euclidean metric if we stretch the state space X onto the radius 2 sphere. This stretching is accomplished using the *Akin transformation* $H_i(x) = 2\sqrt{x_i}$, which emphasizes changes in the use of rare strategies relative to changes in the use of common ones (Akin 1979, 1990; Sandholm et al. 2017). (There is also a dynamic that generates the gradient system for f on X under the usual metric: the so-called *projection dynamic* (Lahkar and Sandholm 2017; Nagurney and Zhang 1997; Sandholm et al. 2017).)

Example 21 Consider evolution in 123 Coordination:

	1	2	3
1	1	0	0
2	0	2	0
3	0	0	3

Figure 2a presents a phase diagram of the replicator dynamic on its natural state space X ,



Evolutionary Game Theory, Fig. 2 The replicator dynamic in 123 Coordination. Colors represent the value of the game's potential function

drawn atop of a contour plot of the potential function $f(x) = \frac{1}{2}((x_1)^2 + 2(x_2)^2 + 3(x_3)^2)$. Evidently, all solution trajectories ascend this function and converge to one of the seven symmetric Nash equilibria, with trajectories from all but a measure zero set of initial conditions converging to one of the three pure equilibria.

Figure 2b presents another phase diagram for the replicator dynamic, this time after the solution trajectories and the potential function have been

transported to the surface of the radius 2 sphere using the Akin transformation. In this case, solutions cross the level sets of the potential function orthogonally, moving in the direction that increases potential most quickly.

Stable Games

A population game F is a *stable game* (Hofbauer and Sandholm 2006a) if

$$(y - x)'(F(y) - F(x)) \leq 0 \quad \text{for all } x, y \in X. \quad (21)$$

If the inequality in (21) always holds strictly, then F is a *strictly stable game*.

If F is smooth, then F is a stable game if and only if it satisfies *self-defeating externalities*:

$$z'DF(x)z \leq 0 \quad \text{for all } z \in TX \text{ and } x \in X, \quad (22)$$

where $DF(x)$ is the derivative of $F : X \rightarrow \mathbf{R}^n$ at x . This condition requires that the improvements in the payoffs of strategies to which revising agents are switching are always exceeded by the improvements in the payoffs of strategies which revising agents are abandoning.

Example 22 The symmetric normal form game A is *symmetric zero-sum* if A is skew-symmetric (i.e., if $A = -A'$), so that the payoffs of the matched players always sum to zero. (An example is provided by the standard Rock-Paper-Scissors game (Example 4).) Under this assumption, $z'Az = 0$ for all $z \in \mathbf{R}^n$; thus, the population game generated by random matching in A , $F(x) = Ax$, is a stable game that is not strictly stable.

Example 23 Suppose that A satisfies the interior ESS condition (18). Then (22) holds strictly, so $F(x) = Ax$ is a strictly stable game. Examples satisfying this condition include the Hawk-Dove game (Example 3) and any good Rock-Paper-Scissors game (Example 4).

Example 24 A *war of attrition* (Bishop and Cannings 1978) is a symmetric normal form

game in which strategies represent amounts of time committed to waiting for a scarce resource. If the two players choose times i and $j > i$, then the j player obtains the resource, worth v , while both players pay a cost of c_i : once the first player leaves, the other seizes the resource immediately. If both players choose time i , the resource is split, so payoffs are $\frac{v}{2} - c_i$ each. It can be shown that for any resource value $v \in \mathbf{R}$ and any increasing cost vector $c \in \mathbf{R}^n$, random matching in a war of attrition generates a stable game (Hofbauer and Sandholm 2006a).

The flavor of the self-defeating externalities condition (22) suggests that obedience of incentives will push the population toward some “central” equilibrium state. In fact, the set of Nash equilibria of a stable game is always convex, and in the case of strictly stable games, equilibrium is unique. Moreover, it can be shown that the replicator dynamic converges to Nash equilibrium from all interior initial conditions in any strictly stable game (Akin 1990; Hofbauer et al. 1979; Zeeman 1980), and that the direct evaluation dynamics introduced above converge to Nash equilibrium from all initial conditions in all stable games, strictly stable or not (Hofbauer 2000; Hofbauer and Sandholm 2006a, 2007; Smith 1984). In each case, the proof of convergence is based on the construction of a Lyapunov function that solutions of the relevant dynamic descend. The Lyapunov functions for the five basic dynamics are presented in Table 3.

Interestingly, the convergence results for direct evaluation dynamics are not restricted to the dynamics listed in Table 3, but extend to other dynamics in the same families (cf Table 2). But compared to the conditions for convergence in potential games, the conditions for convergence in stable games demand additional structure on the adjustment process (Hofbauer and Sandholm 2006a).

Perturbed Best Response Dynamics in Supermodular Games

Supermodular games are defined by the property that higher choices by one’s opponents (with respect to the natural ordering on $S = \{1, \dots, n\}$) make one’s own higher strategies look relatively more

Evolutionary Game Theory, Table 3 Lyapunov functions for five basic deterministic dynamics in stable games

Dynamic	Lyapunov function for stable games
Replicator	$H_{x^*}(x) = \sum_{i \in S(x^*)} x_i^* \log \frac{x_i}{x_i^*}$
Logit	$\tilde{G}(x) = \max_{y \in \text{int}(X)} \left(y' \hat{F}(x) - \eta \sum_{i \in S} y_i \log y_i \right) + \eta \sum_{i \in S} x_i \log x_i$
Best response	$G(x) = \max_{i \in S} \hat{F}_i(x)$
BNN	$\Gamma(x) = \frac{1}{2} \sum_{i \in S} [\hat{F}_i(x)]_+^2$
Smith	$\Psi(x) = \frac{1}{2} \sum_{i \in S} \sum_{j \in S} x_i [F_j(x) - F_i(x)]_+^2$

desirable. Let the matrix $\Sigma \in \mathbf{R}^{(n-1) \times n}$ satisfy $\Sigma_{ij} = 1$ if $j > i$ and $\Sigma_{ij} = 0$ otherwise, so that $\Sigma x \in \mathbf{R}^{n-1}$ is the “decumulative distribution function” corresponding to the “density function” x . The population game F is a *supermodular game* if it exhibits *strategic complementarities*:

$$\begin{aligned} \text{If } \Sigma y \geq \Sigma x, \text{ then } F_{i+1}(y) - F_i(y) \\ \geq F_{i+1}(x) - F_i(x) \text{ for all } i \\ < n \text{ and } x \in X. \end{aligned} \quad (23)$$

If F is smooth, condition (23) is equivalent to

$$\begin{aligned} \frac{\partial(F_{i+1} - F_i)}{\partial(e_{j+1} - e_j)}(x) \geq 0 \text{ for all } i, j \\ < n \text{ and } x \in X. \end{aligned} \quad (24)$$

Example 25 Consider this model of *search with positive externalities*. A population of agents choose levels of search effort in $S = \{1, \dots, n\}$. The payoff to choosing effort i is

$$F_i(x) = m(i)b(a(x)) - c(i),$$

where $a(x) = \sum_{k \leq n} kx_k$ is the aggregate search effort, b is some increasing benefit function, m is an increasing multiplier function, and c is an arbitrary cost function. Notice that the benefits from searching are increasing in both own search effort and in the aggregate search effort. It is easy to check that F is a supermodular game.

Complementarity condition (23) implies that the agents’ best response correspondence is monotone in the stochastic dominance order, which in turn ensures the existence of minimal

and maximal Nash equilibria (Topkis 1979). One can take advantage of the monotonicity of best responses in studying evolutionary dynamics by appealing to the theory of monotone dynamical systems (Smith 1995). To do so, one needs to focus on dynamics that respect the monotonicity of best responses and that also are smooth, so that the theory of monotone dynamics can be applied. It turns out that the logit dynamic satisfies these criteria; so does any perturbed best response dynamic defined in terms of stochastic payoff perturbations. In supermodular games, these dynamics define cooperative differential equations; consequently, solutions of these dynamics from almost every initial condition converge to an approximate Nash equilibrium (Hofbauer and Sandholm 2007).

Imitation Dynamics in Dominance Solvable Games Suppose that in the population game F , strategy i is strictly dominated by strategy j : $F_i(x) < F_j(x)$ for all $x \in X$. Consider the evolution of behavior under the replicator dynamic (13). Since for this dynamic we have

$$\begin{aligned} \frac{d x_i}{d t x_j} &= \frac{\dot{x}_i x_j - \dot{x}_j x_i}{(x_j)^2} \\ &= \frac{x_i \hat{F}_i(x) x_j - x_j \hat{F}_j(x) x_i}{(x_j)^2} \\ &= \frac{x_i}{x_j} (\hat{F}_i(x) - \hat{F}_j(x)), \end{aligned}$$

solutions from every interior initial condition converge to the face of the simplex where the dominated strategy is unplayed (Akin 1980). It follows that the

replicator dynamic converges in games with a strictly dominant strategy, and by iterating this argument, one can show that this dynamic converges to equilibrium in any game that can be solved by iterative deletion of strictly dominated strategies. In fact, this argument is not specific to the replicator dynamic, but can be shown to apply to a range of dynamics based on imitation (Hofbauer and Weibull 1996; Samuelson and Zhang 1992). Even in games which are not dominance solvable, arguments of a similar flavor can be used to restrict the long run behavior of imitative dynamics to better-reply closed sets (Ritzberger and Weibull 1995); see section “Convergence to Equilibria and to Better-Reply Closed Sets” for a related discussion.

While the analysis here has focused on imitative dynamics, it is natural to expect that elimination of dominated strategies will extend to any reasonable evolutionary dynamic. But we will see in section “Survival of Dominated Strategies” that this is not the case: the elimination of dominated strategies that obtains under imitative dynamics is the exception, not the rule.

Nonconvergence

The previous section revealed that when certain global structural conditions on payoffs are satisfied, one can establish global convergence to equilibrium under various classes of evolutionary dynamics. Of course, if these conditions are not met, convergence cannot be guaranteed. In this section, we offer examples to illustrate some of the possibilities for nonconvergent limit behavior.

Conservative Properties of the Replicator Dynamic in Zero-Sum Games

In section “Stable Games,” we noted that in strictly stable games, the replicator dynamic converges to Nash equilibrium from all interior initial conditions. To prove this, one shows that interior solutions descend the function

$$H_{x^*}(x) = \sum_{i \in S(x^*)} x_i^* \log \frac{x_i}{x_i^*},$$

until converging to its minimizer, the unique Nash equilibrium x^* .

Now, random matching in a symmetric zero-sum game generates a population game that is

stable, but not strictly stable (Example 22). In this case, for each interior Nash equilibrium x^* , the function H_{x^*} is a constant of motion for the replicator dynamic: its value is fixed along every interior solution trajectory.

Example 26 Suppose that agents are randomly matched to play the symmetric zero-sum game A , given by

	1	2	3	4
1	0	-1	0	1
2	1	0	-1	0
3	0	1	0	-1
4	-1	0	1	0

The Nash equilibria of $F(x) = Ax$ are the points on the line segment NE connecting states $(\frac{1}{2}, 0, \frac{1}{2}, 0)$ and $(0, \frac{1}{2}, 0, \frac{1}{2})$, a segment that passes through the barycenter $x^* = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. Figure 3 shows solutions to the replicator dynamic that lie on the level set $H_{x^*}(x) = .58$. Evidently, each of these solutions forms a closed orbit.

Although solution trajectories of the replicator dynamic do not converge in zero-sum games, it can be proved that the time average of each solution trajectory converges to Nash equilibrium (Schuster et al. 1981).

The existence of a constant of motion is not the only conservative property enjoyed by replicator dynamics for symmetric zero-sum games: these dynamics are also volume preserving after an appropriate change of speed or change of measure (Akin and Losert 1984; Hofbauer 1995a).


Games with Nonconvergent Dynamics

The conservative properties described in the previous section have been established only for the replicator dynamic (and its distant relative, the projection dynamic (Sandholm et al. 2017)). Inspired by Shapley (1964), many researchers have sought to construct games in which large classes of evolutionary dynamics fail to converge to equilibrium.

Example 27 Suppose that players are randomly matched to play the following symmetric normal form game (Hofbauer and Sigmund 1998; Hofbauer and Swinkels 1996):

Evolutionary Game

Theory, Fig. 3 Solutions of the replicator dynamic in a zero-sum game. The solutions pictured lie on the level set $H_{x^*}(x) = .58$



	1	2	3	4
1	0	0	-1	ε
2	ε	0	0	-1
3	-1	ε	0	0
4	0	-1	ε	0

When $\varepsilon = 0$, the payoff matrix $A^\varepsilon = A^0$ is symmetric, so F^0 is a potential game with potential function $f(x) = \frac{1}{2}x'A^0x = -x_1x_3 - x_2x_4$. The function f attains its minimum of $-\frac{1}{4}$ at states $v = (\frac{1}{2}, 0, \frac{1}{2}, 0)$ and $w = (0, \frac{1}{2}, 0, \frac{1}{2})$, has a saddle point with value $-\frac{1}{8}$ at the Nash equilibrium $x^* = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$, and attains its maximum of 0 along the closed path of Nash equilibria γ consisting of edges e_1e_2 , e_2e_3 , e_3e_4 and e_4e_1 . Let $\dot{x} = V^F(x)$ be an evolutionary dynamic that satisfies Nash stationarity (NS) and positive correlation (PC), and that is based on a revision protocol that is continuous (C). If we apply this dynamic to game F^0 , then the foregoing discussion implies that all solutions to $\dot{x} = V^{F^0}(x)$ whose initial conditions ξ satisfy $f(\xi) > -\frac{1}{8}$ converge to γ . The Smith dynamic for F^0 is illustrated in Fig. 4a.

Now consider the same dynamic for the game F^ε , where $\varepsilon > 0$. By continuity (C), the attractor γ of V^{F^0} continues to an attractor γ^ε of V^{F^ε} whose basin of attraction approximates that of γ under $\dot{x} = V^{F^0}(x)$ (Fig. 4b). But since the unique Nash equilibrium of F^ε is the barycenter x^* , it follows that solutions from most initial conditions converge to an attractor far from any Nash equilibrium.

Other examples of games in which many dynamics fail to converge include monocyclic

games (Benaïm et al. 2006; Gaunersdorfer and Hofbauer 1995; Hofbauer 1995b; Hofbauer and Sigmund 1988), Mismatching Pennies (Hart and Mas-Colell 2003; Jordan 1993), and the hypnodisk game (Hofbauer and Sandholm 2006b). These examples demonstrate that there is no evolutionary dynamic that converges to Nash equilibrium regardless of the game at hand. This suggests that in general, analyses of long run behavior should not restrict attention to equilibria alone.

Chaotic Dynamics

We have seen that deterministic evolutionary game dynamics can follow closed orbits and approach limit cycles. We now show that they also can behave chaotically.


Example 28 Consider evolution under the replicator dynamic when agents are randomly matched to play the symmetric normal form game below (Arneodo et al. 1980; Skyrms 1992), whose lone interior Nash equilibrium is the barycenter $x^* = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$:

	1	2	3	4
1	0	-12	0	22
2	20	0	0	-10
3	-21	-4	0	35
4	10	-2	2	0

Figure 5 presents a solution to the replicator dynamic for this game from initial condition $x_0 = (.24, .26, .25, .25)$. This solution spirals

Evolutionary Game

Theory, Fig. 4 Solutions of the Smith dynamic in (a) the potential game F^0 ; (b) the perturbed potential game F^ε , $\varepsilon = \frac{1}{10}$



clockwise about x^* . Near the rightmost point of each circuit, where the value of x_3 gets close to zero, solutions sometimes proceed along an “outside” path on which the value of x_3 surpasses .6. But they sometimes follow an “inside” path on which x_3 remains below .4, and at other times do something in between. Which of these alternatives occurs is difficult to predict from approximate information about the previous behavior of the system.

While the game in Example 28 has a complicated payoff structure, in multipopulation contexts one can find chaotic evolutionary dynamics in very simple games (Sato et al. 2002).

Survival of Dominated Strategies


In section “Imitation Dynamics in Dominance Solvable Games,” we saw that dynamics based on imitation eliminate strictly dominated strategies along solutions from interior initial conditions. While this result seems unsurprising, it is actually extremely fragile: (Berger and Hofbauer 2006;

Hofbauer and Sandholm 2006b) prove that dynamics that satisfy continuity (C), Nash stationarity (NS), and positive correlation (PC) and that are not based exclusively on imitation must fail to eliminate strictly dominated strategies in some games. Thus, evolutionary support for a basic rationality criterion is more tenuous than the results for imitative dynamics suggest.

Example 29 Figure 6a presents the Smith dynamic for “bad RPS with a twin”:

	R	P	S	T
R	0	-2	1	1
P	1	0	-2	-2
S	-2	1	0	0
T	-2	1	0	0

The Nash equilibria of this game are the states on line segment $NE = \{x^* \in X : x^* = (\frac{1}{3}, \frac{1}{3}, c, \frac{1}{3} - c)\}$, which is a repeller under the Smith dynamic. Under



Evolutionary Game Theory, Fig. 5 Chaotic behavior under the replicator dynamic

this dynamic, strategies gain players at rates that depend on their payoffs, but lose players at rates proportional to their current usage levels. It follows that when the dynamics are not at rest, the proportions of players choosing strategies 3 and 4 become equal, so that the dynamic approaches the plane $P = \{x \in X : x_3 = x_4\}$ on which the twins receive equal weight. Since the usual three-strategy version of bad RPS, exhibits cycling solutions here on the plane P approach a closed orbit away from any Nash equilibrium.


Figure 6b presents the Smith dynamic in “bad RPS with a feeble twin,”

	R	P	S	T
R	0	-2	1	1
P	1	0	-2	-2
S	-2	1	0	0
T	$-2 - \epsilon$	$1 - \epsilon$	$-\epsilon$	$-\epsilon$

with $\epsilon = \frac{1}{10}$. Evidently, the attractor from Fig. 6a moves slightly to the left, reflecting the fact that the payoff to Twin has gone down. But since the new attractor is in the interior of X , the strictly dominated strategy Twin is always played with probabilities bounded far away from zero.

Stochastic Dynamics

In section “Revision Protocols” we defined the stochastic evolutionary process $\{X_t^N\}$ in terms of



Evolutionary Game Theory, Fig. 6 The Smith dynamic in two games

a simple model of myopic individual choice. We then turned to the study of deterministic dynamics, which we claimed could be used to approximate the stochastic process $\{X_t^N\}$ over finite time spans and for large population sizes. In this section, we turn our attention to the stochastic process $\{X_t^N\}$ itself. After offering a formal version of the deterministic approximation result, we investigate the long run behavior of $\{X_t^N\}$, focusing on the questions of convergence to equilibrium and selection among multiple stable equilibria.

Deterministic Approximation

In section “Revision Protocols,” we defined the Markovian evolutionary process $\{X_t^N\}$ from a revision protocol ρ , a population game F , and a finite population size N . In section “Mean

Dynamics,” we argued that the expected motion of this process is captured by the mean dynamic


$$\begin{aligned}\dot{x}_i &= V_i^F(x) \\ &= \sum_{j \in S} x_j \rho_{ji}(F(x), x) \\ &\quad - x_i \sum_{j \in S} \rho_{ij}(F(x), x).\end{aligned}\quad (\text{M})$$

The basic link between the Markov process $\{X_t^N\}$ and its mean dynamic (M) is provided by *Kurtz’s Theorem* (Kurtz 1970), variations and extensions of which have been offered in a number of game-theoretic contexts (Benaïm and Weibull 2003; Binmore and Samuelson 1997; Börgers and Sarin 1997; Boylan 1995; Sandholm 2003; Tanabe 2006). Consider the sequence of Markov processes $\left\{ \left\{ X_t^N \right\}_{t \geq 0} \right\}_{N=N_0}^{\infty}$, supposing that the initial conditions X_0^N converge to $x_0 \in X$. Let $\{x_t\}_{t \geq 0}$ be the solution to the mean dynamic (M) starting from x_0 . Kurtz’s Theorem tells us that for each finite time horizon $T < \infty$ and error bound $\varepsilon > 0$, we have that

$$\lim_{N \rightarrow \infty} P \left(\sup_{t \in [0, T]} |X_t^N - x_t| < \varepsilon \right) = 1. \quad (25)$$

Thus, when the population size N is large, nearly all sample paths of the Markov process $\{X_t^N\}$ stay within ε of a solution of the mean dynamic (M) through time T . By choosing N large enough, we can ensure that with probability close to one, X_t^N and x_t differ by no more than ε for all times t between 0 and T (Fig. 7).

The intuition for this result comes from the law of large numbers. At each revision opportunity, the increment in the process $\{X_t^N\}$ is stochastic. Still, at most population states the expected number of revision opportunities that arrive during the brief time interval $I = [t, t + dt]$ is large – in particular, of order Ndt . Since each opportunity leads to an increment of the state of size $\frac{1}{N}$, the size of the overall change in the state during time interval I is of order dt . Thus, during this interval there are a large number of revision opportunities, each following nearly the same transition probabilities, and hence having nearly



Evolutionary Game Theory, Fig. 7 Deterministic approximation of the Markov process $\{X_t^N\}$

the same expected increments. The law of large numbers therefore suggests that the change in $\{X_t^N\}$ during this interval should be almost completely determined by the expected motion of $\{X_t^N\}$, as described by the mean dynamic (M).

Convergence to Equilibria and to Better-Reply Closed Sets

Stochastic models of evolution can also be used to address directly the question of convergence to equilibrium (Dindoš and Mezzetti 2006; Friedman and Mezzetti 2001; Josephson 2008; Josephson and Matros 2004; Kukushkin 2004; Monderer and Shapley 1996; Sandholm 2001a; Peyton Young 1993a). Suppose that a society of agents is randomly matched to play an (asymmetric) normal form game that is *weakly acyclic in better replies*: from each strategy profile, there exists a sequence of profitable unilateral deviations leading to a Nash equilibrium. If agents switch to strategies that do at least as well as their current one against the choices of random samples of opponents, then the society will eventually escape any better-response cycle, ultimately settling upon a Nash equilibrium.

Importantly, many classes of normal form games are weakly acyclic in better replies: these include potential games, dominance solvable games, certain supermodular games, and certain *aggregative games*, in which each agent’s payoffs only depend on opponents’ behavior through a

scalar aggregate statistic. Thus, in all of these cases, simple stochastic better-reply procedures are certain to lead to Nash equilibrium play.

Outside these classes of games, one can narrow down the possibilities for long run behavior by looking at *better-reply closed sets*: that is, subsets of the set of strategy profiles that cannot be escaped without a player switching to an inferior strategy (cf. Basu and Weibull 1991; Ritzberger and Weibull 1995). Stochastic better-reply procedures must lead to a cluster of population states corresponding to a better-reply closed set; once the society enters such a cluster, it never departs.

Stochastic Stability and Equilibrium Selection

To this point, we used stochastic evolutionary dynamics to provide foundations for deterministic dynamics and to address the question of convergence to equilibrium. But stochastic evolutionary dynamics introduce an entirely new possibility: that of obtaining unique long-run predictions of play, even in games with multiple locally stable equilibria. This form of analysis, which we consider next, was pioneered by Foster and Peyton Young (1990), Kandori et al. (1993), and Peyton Young (1993a), building on mathematical techniques due to Freidlin and Wentzell (1998).

Stochastic Stability

To minimize notation, let us describe the evolution of behavior using a discrete-time Markov chain $\{X_k^{N,\varepsilon}\}_{k=0}^{\infty}$ on $\mathcal{X}^N$, where the parameter $\varepsilon > 0$ represents the level of “noise” in agents’ decision procedures. The noise ensures that the Markov chain is irreducible and aperiodic: any state in $\mathcal{X}^N$ can be reached from any other, and there is positive probability that a period passes without a change in the state.

Under these conditions, the Markov chain $\{X_k^{N,\varepsilon}\}$ admits a unique *stationary distribution*, $\mu^{N,\varepsilon}$, a measure on the state space $\mathcal{X}^N$ that is invariant under the Markov chain:

$$\sum_{x \in \mathcal{X}^N} \mu^{N,\varepsilon}(x) P(X_{k+1}^{N,\varepsilon} = y | X_k^{N,\varepsilon} = x) = \mu^{N,\varepsilon}(y) \quad \text{for all } y \in \mathcal{X}^N.$$

The stationary distribution describes the long run behavior of the process $\{X_t^{N,\varepsilon}\}$ in two distinct ways. First, $\mu^{N,\varepsilon}$ is the *limiting distribution* of $\{X_t^{N,\varepsilon}\}$:

$$\lim_{k \rightarrow \infty} P(X_k^{N,\varepsilon} = y | X_0^{N,\varepsilon} = x) = \mu^{N,\varepsilon}(y) \quad \text{for all } x, y \in \mathcal{X}^N.$$

Second, $\mu^{N,\varepsilon}$ almost surely describes the *limiting empirical distribution* of $\{X_t^{N,\varepsilon}\}$:

$$P\left(\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} 1_{\{X_k^{N,\varepsilon} \in A\}} = \mu^{N,\varepsilon}(A)\right) = 1 \quad \text{for any } A \subseteq \mathcal{X}^N.$$

Thus, if most of the mass in the stationary distribution $\mu^{N,\varepsilon}$ were placed on a single state, then this state would provide a unique prediction of long run behavior.

With this motivation, consider a sequence of Markov chains $\left\{\left\{X_k^{N,\varepsilon}\right\}_{k=0}^{\infty}\right\}_{\varepsilon \in (0,\bar{\varepsilon})}$ parametrized by noise levels ε that approach zero. Population state $x \in \mathcal{X}^N$ is said to be *stochastically stable* if it retains positive weight in the stationary distributions of these Markov chains as ε becomes arbitrarily small:

$$\lim_{\varepsilon \rightarrow 0} \mu^{N,\varepsilon}(x) > 0.$$

When the stochastically stable state is unique, it offers a unique prediction of play that is relevant over sufficiently long time spans.

Bernoulli Arrivals and Mutations

Following the approach of many early contributors to the literature, let us consider a model of stochastic evolution based on *Bernoulli arrivals of revision opportunities* and *best responses with mutations*. The former assumption means that during each discrete time period, each agent has probability $\theta \in (0, 1]$ of receiving an opportunity to update his strategy. This assumption differs than the one we proposed in section “[Revision Protocols](#)”; the key new implication is that all agents may receive revision opportunities simultaneously. (Models that assume this directly generate similar results.) The latter assumption posits

that when an agent receives a revision opportunity, he plays a best response to the current strategy distribution with probability $1 - \varepsilon$, and chooses a strategy at random with probability ε .

Example 30 Suppose that a population of N agents is randomly matched to play the Stag Hunt game (Example 2):

	H	S
H	h	h
S	0	s

Since $s > h > 0$, hunting hare and hunting stag are both symmetric pure equilibria; the game also admits the symmetric mixed equilibrium $x^* = (x_H^*, x_S^*) = (\frac{s-h}{s}, \frac{h}{s})$.

If more than fraction x_H^* of the agents hunt hare, then hare is the unique best response, while if more than fraction x_S^* of the agents hunt stag, then stag is the unique best response. Thus, under any deterministic dynamic that respects payoffs, the mixed equilibrium x^* divides the state space into two basins of attraction, one for each of the two pure equilibria.

Now consider our stochastic evolutionary process. If the noise level ε is small, this process typically behaves like a deterministic process, moving quickly toward one of the two pure states, $e_H = (1, 0)$ or $e_S = (0, 1)$, and remaining there for some time. But since the process is ergodic, it will eventually leave the pure state it reaches first, and in fact will switch from one pure state to the other infinitely often.


To determine the stochastically stable state, we must compute and compare the “improbabilities” of these transitions. If the current state is e_H , a transition to e_S requires mutations to cause roughly Nx_S^* agents to switch to the suboptimal strategy S, sending the population into the basin of attraction of e_S ; the probability of this event is of order $\varepsilon^{Nx_S^*}$. Similarly, to transit from e_S to e_H , mutations must cause roughly $Nx_H^* = N(1 - x_S^*)$ to switch from S to H; this probability of this event is of order $\varepsilon^{N(1-x_S^*)}$.

Which of these rare events is more likely ones depends on whether x_S^* is greater than or less than $\frac{1}{2}$. If $s > 2h$, so that $x_S^* < \frac{1}{2}$, then $\varepsilon^{Nx_S^*}$ is much smaller than $\varepsilon^{N(1-x_S^*)}$, when ε is small; thus, state e_S is stochastically stable (Fig. 8a). If instead $s < 2h$, so that $x_S^* > \frac{1}{2}$, then $\varepsilon^{N(1-x_S^*)} < \varepsilon^{Nx_S^*}$ so e_H is stochastically stable (Fig. 8b).

These calculations show that *risk dominance* – being the optimal response against a uniformly randomizing opponent – drives stochastic stability 2×2 games. In particular, when $s < 2h$, so that risk dominance and payoff dominance disagree, stochastic stability favors the former over the latter.

This example illustrates how under Bernoulli arrivals and mutations, stochastic stability analysis is based on *mutation counting*: that is, on determining how many simultaneous mutations are required to move from each equilibrium into the basin of attraction of each other equilibrium. In games with more than two strategies, completing the argument becomes more complicated than in the example above: the analysis, typically based on the tree-analysis techniques of Freidlin

Evolutionary Game Theory, Fig. 8 Equilibrium selection via mutation counting in Stag Hunt games



and Wentzell (1998) and Peyton Young (1993a), requires one to account for the relative difficulties of transitions between all pairs of equilibria. Ellison (2000) develops a streamlined method of computing the stochastically stable state based on radius-coradius calculations; while this approach is not always sufficiently fine to yield a complete analysis, in the cases where it works it can be considerably simpler to apply than the tree-analysis method.

These techniques have been employed successfully to variety of classes of games, including pure coordination games, supermodular games, games satisfying “bandwagon” properties, and games with equilibria that satisfy generalizations of risk dominance (Ellison 2000; Kandori and Rob 1995, 1998; Maruta 1997). A closely related literature uses stochastic stability as a basis for evaluating traditional solution concepts for extensive form games (Hart 2002; Jacobsen et al. 2001; Kim and Sobel 1995; Kuzmics 2004; Nöldeke and Samuelson 1993; Samuelson 1994, 1997).

A number of authors have shown that variations on the Bernoulli arrivals and mutations model can lead to different equilibrium selection results. For instance, Robson and Vega-Redondo (1996) and Vega-Redondo (1996) show that if choices are determined from the payoffs from a single round of matching (rather than from expected payoffs), the payoff dominant equilibrium rather than the risk dominant equilibrium is selected. If choices depend on strategies’ relative performances rather than their absolute performances, then long run behavior need not resemble a Nash equilibrium at all (Bergin and Bernhardt 2004; Rhode and Stegeman 1996; Sandholm 1998; Stegeman and Rhode 2004). Finally, if the probability of mutation depends on the current population state, then any recurrent set of the unperturbed process (e. g., any pure equilibrium of a coordination game) can be selected in the long run if the mutation rates are specified in an appropriate way (Bergin and Lipman 1996). This last result suggests that mistake probabilities should be provided with an explicit foundation, a topic we take up in section “Poisson Arrivals and Payoff Noise.”

Another important criticism of the stochastic stability literature concerns the length of time needed for its predictions to become relevant (Binmore et al. 1995b; Ellison 1993). If the population size N is large and the mutation rate ε is small, then the probability ε^{cN} that a transition between equilibria occurs during given period is miniscule; the waiting time between transitions is thus enormous. Indeed, if the mutation rate falls over time, or if the population size grows over time, then ergodicity may fail, abrogating equilibrium selection entirely (Robles 1998; Sandholm and Pauzner 1998). These analyses suggest that except in applications with very long time horizons, the unique predictions generated by analyses of stochastic stability may be inappropriate, and that modelers would do better to focus on history-dependent predictions of the sort provided by deterministic models. At the same time, there are frameworks in which stochastic stability becomes relevant much more quickly. The most important of these are local interaction models, which we discuss in section “Local Interaction.”

Poisson Arrivals and Payoff Noise

Combining the assumption of Bernoulli arrivals of revision opportunities with that of best responses with mutations creates a model in which the probabilities of transitions between equilibria are easy to compute: one can focus on events in which large numbers of agents switch to a suboptimal strategy at once, each doing so with the same probability. But the simplicity of this argument also highlights the potency of the assumptions behind it.

An appealing alternative approach is to model stochastic evolution using *Poisson arrivals of revision opportunities* and *payoff noise* (Binmore and Samuelson 1997; Binmore et al. 1995b; Blume 1997, 2003; Dokumaci and Sandholm 2007b; Maruta 2002; Myatt and Wallace 2003; Ui 1998; van Damme and Weibull 2002; Peyton Young 1998b). (One can achieve similar effects by looking at models defined in terms of stochastic differential equations; see (Beggs 2002; Cabrales 2000; Foster and Peyton Young 1990; Fudenberg and Harris 1992; Imhof 2005).) By

allowing revision opportunities to arrive in continuous time, as we did in section “[Revision Protocols](#),” we ensure that agents do not receive opportunities simultaneously, ruling out the simultaneous mass revisions that drive the Bernoulli arrival model. (One can accomplish the same end using a discrete time model by assuming that one agent updates during each period; the resulting process is a random time change away from the Poisson arrivals model.)

Under Poisson arrivals, transitions between equilibria occur gradually, as the population works its way out of basins of attraction one agent at a time. In this context, the mutation assumption becomes particularly potent, ensuring that the probabilities of suboptimal choices do not vary with their payoff consequences. Under the alternative assumption of payoff noise, one supposes that agents play best responses to payoffs that are subject to random perturbations drawn from a fixed multivariate distribution. In this case, suboptimal choices are much more likely near basin boundaries, where the payoffs of second-best strategies are not much less than those of optimal ones, than they are at stable equilibria, where payoff differences are larger.

Evidently, assuming Poisson arrivals and payoff noise means that stochastic stability cannot be assessed by way of mutation counting. To determine the unlikelihood of escaping from an equilibrium’s basin of attraction, one must not only account for the “width” of the basin of attraction (i.e., the number of suboptimal choices needed to escape it), but also for its “depth” (the unlikelihood of each of these choices). In two-strategy games this is not difficult to accomplish: in this case the evolutionary process is a birth-and-death chain, and its stationary distribution can be expressed using an explicit formula. Beyond this case, one can employ the Freidlin and Wentzell (1998) machinery, although doing so tends to be computationally demanding.

This computational burden is less in models that retain Poisson arrivals, but replace perturbed optimization with decision rules based on imitation and mutation (Fudenberg and Imhof 2006). Because agents imitate successful opponents, the population spends the vast majority of periods on

the edges of the simplex, implying that the probabilities of transitions between vertices can be determined using birth-and-death chain methods (Nowak et al. 2004). As a consequence, one can reduce the problem of finding the stochastically stable state in an n strategy coordination game to that of computing the limiting stationary distribution of an n state Markov chain.

Stochastic Stability via Large Population Limits

The approach to stochastic stability followed thus far relies on small noise limits: that is, on evaluating the limit of the stationary distributions $\mu^{N, \varepsilon}$ as the noise level ε approaches zero. Binmore and Samuelson (1997) argue that in the contexts where evolutionary models are appropriate, the amount of noise in agents decisions is not negligible, so that taking the low noise limit may not be desirable. At the same time, evolutionary models are intended to describe behavior in large populations, suggesting an alternative approach: that of evaluating the limit of the stationary distributions $\mu^{N, \varepsilon}$ as the population size N grows large.

In one respect, this approach complicates the analysis. When N is fixed and ε varies, each stationary distribution $\mu^{N, \varepsilon}$ is a measure on the fixed state space $\mathcal{X}^N = \{x \in X : Nx \in \mathbf{Z}^n\}$. But when ε is fixed and N varies, the state space $\mathcal{X}^N$ varies as well, and one must introduce notions of weak convergence of probability measures in order to define stochastic stability.

But in other respects taking large population limits can make analysis simpler. We saw in section “[Deterministic Approximation](#)” that by taking the large population limit, we can approximate the finite-horizon sample paths of the stochastic evolutionary process $\{X_t^{N, \varepsilon}\}$ by solutions to the mean dynamic (M). Now we are concerned with infinite horizon behavior, but it is still reasonable to hope that the large population limit will again reduce some of our computations to a calculus problems.

As one might expect, this approach is easiest to follow in the two-strategy case, where for each fixed population size N , the evolutionary process $\{X_t^{N, \varepsilon}\}$ is a birth-and-death chain. When one

takes the large population limit, the formulas for waiting times and for the stationary distribution can be evaluated using integral approximations (Benaïm and Weibull 2003; Binmore and Samuelson 1997; Blume 2003; Peyton Young 1998b). Indeed, the approximations so obtained take an appealing simple form (Sandholm 2007d).

The analysis becomes more complicated beyond the two-strategy case, but certain models have proved amenable to analysis. For instance Fudenberg and Imhof (2006), characterizes large population stochastic stability in models based on imitation and mutation. Imitation ensures that the population spends nearly all periods on the edges of the simplex X , and the large population limit makes evaluating the probabilities of transitions along these edges relatively simple.

If one supposes that agents play best responses to noisy payoffs, then one must account directly for the behavior of the process $\{X_t^{N,\varepsilon}\}$ in the interior of the simplex. One possibility is to combine the deterministic approximation results from section “[Deterministic Approximation](#)” with techniques from the theory of stochastic approximation (Benaïm 1998; Benaïm and Hirsch 1999) to show that the large N limiting stationary distribution is concentrated on attractors of the mean dynamic. By combining this idea with convergence results for deterministic dynamics from section “[Global Convergence](#),” Hofbauer and Sandholm (2007) shows that the limiting stationary distribution must be concentrated around equilibrium states in potential games, stable games, and supermodular games.

The results in Hofbauer and Sandholm (2007) do not address the question of equilibrium selection. However, for the specific case of logit evolution in potential games, a complete characterization of the large population limit of the process $\{X_t^{N,\varepsilon}\}$ has been obtained (Benaïm and Sandholm 2007). By combining deterministic approximation results, which describe the usual behavior of the process within basins of attraction, with a large deviations analysis, which characterizes the rare escapes from basins of attraction, one can obtain a precise asymptotic formula for the large N limiting stationary distribution. This

formula accounts both for the typical procession of the process along solutions of the mean dynamic, and for the rare sojourns of the process against this deterministic flow.

Local Interaction

All of the game dynamics considered so far have been based implicitly on the assumption of *global interaction*: each agent’s payoffs depend directly on all agents’ actions. In many contexts, one expects to the contrary that interactions will be local in nature: for instance, agents may live in fixed locations and interact only with neighbors. In addition to providing a natural fit for these applications, *local interaction* models respond to some of the criticisms of the stochastic stability literature. At the same time, once one moves beyond relatively simple cases, local interaction models become exceptionally complicated, and so lend themselves to methods of analysis very different from those considered thus far.

Stochastic Stability and Equilibrium Selection Revisited

In section “[Stochastic Stability and Equilibrium Selection](#),” we saw the prediction of risk dominant equilibrium play provided by stochastic stability models is subverted by the waiting-time critique: namely, that the length of time required before this equilibrium is reached may be extremely long. Ellison (1993, 2000) shows that if interactions are local, then selection of the risk dominant equilibrium persists, and waiting times are no longer an issue.

Example 31 In the simplest local interaction model, a population of N agents are located at N distinct positions around a circle. During each period of play, each agent plays the Stag Hunt game (Examples 2 and 30) with his two nearest neighbors, following the same action against both of his opponents. If we suppose that $s \in (h, 2h)$, so that hunting hare is the risk dominant strategy, then by definition, an agent whose neighbors play different strategies finds it optimal to choose H himself.

Now suppose that there are Bernoulli arrivals of revision opportunities, and that decisions are based on best responses and rare mutations. To move from the all S state to the all H state, it is enough that a single agent mutates S to H . This one mutation begins a chain reaction: the mutating agent's neighbors respond optimally by switching to H themselves; they are followed in this by their own neighbors; and the contagion continues until all agents choose H . Since a single mutation is always enough to spur the transition from all S to all H , the expected wait before this transition is small, even when the population is large.

In contrast, the transition back from all H to all S is extremely unlikely. Even if all but one of the agents simultaneously mutate to S , the contagion process described above will return the population to the all- H state. Thus, while the transition from all- S to all- H occurs quickly, the reverse transition takes even longer than in the global interaction setting.

The local interaction approach to equilibrium selection has been advanced in a variety of directions: by allowing agents to choose their locations (Ely 2002), or to pay a cost to choose different strategies against different opponents (Goyal and Janssen 1997), and by basing agents' decisions on the attainment of aspiration levels (Anderlini and Ianni 1996), or on imitation of successful opponents (Alós-Ferrer and Weidenholzer 2006a, b). A portion of this literature initiated by Blume develops connections between local interaction models in evolutionary game theory with models from statistical mechanics (Blume 1993, 1995, 1997; Kosfeld 2002; Miękisz 2004). These models provide a point of departure for research on complex spatial dynamics in games, which we consider next.

Complex Spatial Dynamics

The local interaction models described above address the questions of convergence to equilibrium and selection among multiple equilibria. In the cases where convergence and selection results obtain, behavior in these models is relatively simple, as most periods are spent with most agents coordinating on a single strategy. A distinct branch of the

literature on evolution and local interaction focuses on cases with complex dynamics, where instead of settling quickly into a homogeneous, static configuration, behavior remains in flux, with multiple strategies coexisting for long periods of time.

Example 32 Cooperating is a dominated strategy in the Prisoner's Dilemma, and is not played in equilibrium in finitely repeated versions of this game. Nevertheless, a pair of Prisoner's Dilemma tournaments conducted by Axelrod (1984) were won by the strategy Tit-for-Tat, which cooperates against cooperative opponents and defects against defectors. Axelrod's work spawned a vast literature aiming to understand the persistence of individually irrational but socially beneficial behavior.

To address this question, Nowak and May (Nowak 2006; Nowak et al. 1994a, b; Nowak and May 1992, 1993) consider a population of agents who are repeatedly matched to play the Prisoner's Dilemma

	C	D
C	1	$-\varepsilon$
D	g	0

where the greedy payoff g exceeds 1 and $\varepsilon > 0$ is small. The agents are positioned on a two-dimensional grid. During each period, each agent plays the Prisoner's Dilemma with the eight agents in his (Moore) neighborhood. In the simplest version of the model, all agents simultaneously update their strategies at the end of each period. If an agent's total payoff that period is as high as that of any of neighbor, he continues to play the same strategy; otherwise, he switches to the strategy of the neighbor who obtained the highest payoff.


Since defecting is a dominant strategy in the Prisoner's Dilemma, one might expect the local interaction process to converge to a state at which all agents defect, as would be the case in nearly any model of global interaction. But while an agent is always better off defecting himself, he also is better off the more of his neighbors cooperate; and since evolution is based on imitation, cooperators tend to have more cooperators as neighbors than do defectors.

In Figs. 9, 10, and 11, we present snapshots of the local interaction process for choices of the greedy payoff g from each of three distinct parameter regions. If $g > \frac{5}{3}$ (Fig. 9), the process quickly converges to a configuration containing a few rectangular islands of cooperators in a sea of defectors; the exact configuration depending on the initial conditions. If instead $g < \frac{8}{5}$ (Fig. 10), the process moves towards a configuration in which agents other than those in a “web” of defectors cooperate. But for $g \in (\frac{8}{5}, \frac{5}{3})$ (Fig. 11), the system evolves in a complicated fashion, with clusters of cooperators and of defectors forming, expanding, disappearing, and reforming. But while the configuration of behavior never stabilizes, the proportion of cooperators appears to settle down to about .30. The specification of the dynamics considered above, based on simultaneous updating and certain imitation of the most successful neighbor, presents a relatively favorable environment for

cooperative behavior. Nevertheless, under Poisson arrivals of revision opportunities, or probabilistic decision rules, or both, cooperation can persist for very long periods of time for values of g significantly larger than 1 (Nowak et al. 1994a, b).


The literature on complex spatial dynamics in evolutionary game models is large and rapidly growing, with the evolution of behavior in the spatial Prisoners’ Dilemma being the single most-studied environment. While analyses are typically based on simulations, analytical results have been obtained in some relatively simple settings (Eshel et al. 1998; Herz 1994).

Recent work on complex spatial dynamics has considered games with three or more strategies, including Rock- Paper-Scissors games, as well as public good contribution games and Prisoner’s Dilemmas with voluntary participation. Introducing more than two strategies can lead to qualitatively novel dynamic phenomena, including large-scale




Evolutionary Game Theory, Fig. 9 Local interaction in a Prisoner’s Dilemma; greedy payoff $g = 1.7$. In Figs. 9, 10, and 11, agents are arrayed on a 100×100 grid with periodic boundaries (i.e., a torus). Initial conditions are random with 75% cooperators and 25% defectors. Agents update simultaneously, imitating the neighbor who earned

the highest payoff. *Blue* cells represent cooperators who also cooperated last period, *green* cells represent new cooperators; *red* cells represent defectors who also defected last period, *yellow* cells represent new defectors. (Figs. 9, 10, and 11 created using VirtualLabs (Hauert 2007))



Evolutionary Game Theory, Fig. 10 Local interaction in a Prisoner’s Dilemma; greedy payoff $g = 1.55$



Evolutionary Game Theory, Fig. 11 Local interaction in a Prisoner's Dilemma; greedy payoff $g = 1.65$

spatial cycles and traveling waves (Hauert et al. 2002; Szabó and Hauert 2002; Tainaka 2001). In addition to simulations, the analysis of complex spatial dynamics is often based on approximation techniques from non-equilibrium statistical physics, and much of the research on these dynamics has appeared in the physics literature. Szabó and Fáth (2007) offers a comprehensive survey of work on this topic.

Applications

Evolutionary game theory was created with biological applications squarely in mind. In the pre-history of the field, Fisher (1930) and Hamilton (1967) used game-theoretic ideas to understand the evolution of sex ratios. Maynard Smith (1972, 1974, 1982; Maynard Smith and Price 1973) introduced his definition of ESS as a way

of understanding ritualized animal conflicts. Since these early contributions, evolutionary game theory has been used to study a diverse array of biological questions, including mate choice, parental investment, parent-offspring conflict, social foraging, and predator-prey systems. For overviews of research on these and other topics in biology, see (Dugatkin and Reeve 1998; Hammerstein and Selten 1994).

The early development of evolutionary game theory in economics was motivated primarily by theoretical concerns: the justification of traditional game-theoretic solution concepts, and the development of methods for equilibrium selection in games with multiple stable equilibria. More recently, evolutionary game theory has been applied to concrete economic environments, in some instances as a means of contending with equilibrium selection problems, and in others to obtain an explicitly dynamic model of the

phenomena of interest. Of course, these applications are most successful when the behavioral assumptions that underlie the evolutionary approach are appropriate, and when the time horizon needed for the results to become relevant corresponds to the one germane to the application at hand.

Topics in economics theoretical studied using the methods of evolutionary game theory range from behavior in markets (Agastya 2004; Alós-Ferrer 2005; Alós-Ferrer et al. 2000, 2006; Ania et al. 2002; Ben-Shoham et al. 2004; Droste et al. 2002; Hopkins and Seymour 2002; Lahkar 2007; Vega-Redondo 1997), to bargaining and hold-up problems (Binmore et al. 2003; Burke and Peyton Young 2001; Dawid and Bentley MacLeod 2008; Ellingsen and Robles 2002; Robles 2008; Tröger 2002; Peyton Young 1993b, 1998a, b), to externality and implementation problems (Cabrales 1999; Cabrales and Ponti 2000; Mathevet 2007; Sandholm 2002, 2005b, 2007b), to questions of public good provision and collective action (Myatt and Wallace 2007, 2008a, b). The techniques described here are being applied with increasing frequency to problems of broader social science interest, including residential segregation (Bøg 2006; Dokumaci and Sandholm 2007a; Möbius 2000; Peyton Young 1998b, 2001; Zhang 2004a, b) and cultural evolution (Bisin and Verdier 2001; Kuran and Sandholm 2008), and to the study of behavior in transportation and computer networks (Fischer and Vöcking 2006; Monderer and Shapley 1996; Nagurney and Zhang 1997; Sandholm 2001b, 2003, 2005b; Smith 1984). A proliferating branch of research extends the approaches described in this article to address the evolution of structure and behavior in social networks; a number of recent books (Goyal 2007; Jackson 2017; Vega-Redondo 2007) offer detailed treatments of work in this domain.

Future Directions

Evolutionary game theory is a maturing field; many basic theoretical issues are well understood, but many difficult questions remain. It is tempting to say that stochastic and local interaction models

offer the more open terrain for further explorations. But while it is true that we know less about these models than about deterministic evolutionary dynamics, even our knowledge of the latter is limited: while dynamics on one and two dimensional state spaces, and for games satisfying a few interesting structural assumptions, are well-understood, the dynamics of behavior in the vast majority of many-strategy games are not.

The prospects for further applications of the tools of evolutionary game theory are brighter still. In economics, and in other social sciences, the analysis of mathematical models has too often been synonymous with the computation and evaluation of equilibrium behavior. The questions of whether and how equilibrium will come to be are often ignored, and the possibility of long-term disequilibrium behavior left unmentioned. For settings in which its assumptions are tenable, evolutionary game theory offers a host of techniques for modeling the dynamics of economic behavior. The exploitation of the possibilities for a deeper understanding of human social interactions has hardly begun.

Acknowledgments The figures in sections “*Deterministic Dynamics*” and “*Local Interaction*” were created using *Dynamo* (Sandholm and Dokumaci 2007) and *VirtualLabs* (Hauert 2007), respectively. I am grateful to Caltech for its hospitality as I completed this article, and I gratefully acknowledge financial support under NSF Grant SES-0617753.

Bibliography

- Agastya M (2004) Stochastic stability in a double auction. *Games Econ Behav* 48:203–222
- Akin E (1979) *The geometry of population genetics*. Springer, Berlin
- Akin E (1980) Domination or equilibrium. *Math Biosci* 50:239–250
- Akin E (1990) The differential geometry of population genetics and evolutionary games. In: Lessard S (ed) *Mathematical and statistical developments of evolutionary theory*. Kluwer, Dordrecht, pp 1–93
- Akin E, Losert V (1984) Evolutionary dynamics of zero-sum games. *J Math Biol* 20:231–258
- Alós-Ferrer C (2005) The evolutionary stability of perfectly competitive behavior. *Econ Theory* 26:497–516
- Alós-Ferrer C, Weidenholzer S (2006a) Contagion and efficiency. *J Econ Theory*, University of Konstanz and University of Vienna

- Alós-Ferrer C, Weidenholzer S (2006b) Imitation, local interactions, and efficiency. *Econ Lett* 93:163–168
- Alós-Ferrer C, Ania AB, Schenk-Hoppé KR (2000) An evolutionary model of Bertrand oligopoly. *Games Econ Behav* 33:1–19
- Alós-Ferrer C, Kirchsteiger G, Walzl M (2006) On the evolution of market institutions: the platform design paradox. Unpublished manuscript, University of Konstanz
- Anderlini L, Ianni A (1996) Path dependence and learning from neighbors. *Games Econ Behav* 13:141–177
- Ania AB, Tröger T, Wambach A (2002) An evolutionary analysis of insurance markets with adverse selection. *Games Econ Behav* 40:153–184
- Arneodo A, Couillet P, Tresser C (1980) Occurrence of strange attractors in three-dimensional Volterra equations. *Phys Lett* 79A:259–263
- Axelrod R (1984) The evolution of cooperation. Basic Books, New York
- Balkenborg D, Schlag KH (2001) Evolutionarily stable sets. *Int J Game Theory* 29:571–595
- Basu K, Weibull JW (1991) Strategy sets closed under rational behavior. *Econ Lett* 36:141–146
- Beckmann M, McGuire CB, Winsten CB (1956) Studies in the economics of transportation. Yale University Press, New Haven
- Beggs AW (2002) Stochastic evolution with slow learning. *Econ Theory* 19:379–405
- Ben-Shoham A, Serrano R, Volij O (2004) The evolution of exchange. *J Econ Theory* 114:310–328
- Benaïm M (1998) Recursive algorithms, urn processes, and the chaining number of chain recurrent sets. *Ergod Theory Dyn Syst* 18:53–87
- Benaïm M, Hirsch MW (1999) On stochastic approximation algorithms with constant step size whose average is cooperative. *Ann Appl Probab* 30:850–869
- Benaïm M, Sandholm WH (2007) Logit evolution in potential games: reversibility, rates of convergence, large deviations, and equilibrium selection. Unpublished manuscript, Université de Neuchâtel and University of Wisconsin
- Benaïm M, Weibull JW (2003) Deterministic approximation of stochastic evolution in games. *Econometrica* 71:873–903
- Benaïm M, Hofbauer J, Hopkins E (2006) Learning in games with unstable equilibria. Unpublished manuscript, Université de Neuchâtel, University of Vienna and University of Edinburgh
- Berger U, Hofbauer J (2006) Irrational behavior in the Brown-von Neumann-Nash dynamics. *Games Econ Behav* 56:1–6
- Bergin J, Bernhardt D (2004) Comparative learning dynamics. *Int Econ Rev* 45:431–465
- Bergin J, Lipman BL (1996) Evolution with state-dependent mutations. *Econometrica* 64:943–956
- Binmore K, Samuelson L (1997) Muddling through: noisy equilibrium selection. *J Econ Theory* 74:235–265
- Binmore K, Samuelson L (1999) Evolutionary drift and equilibrium selection. *Rev Econ Stud* 66:363–393
- Binmore K, Gale J, Samuelson L (1995a) Learning to be imperfect: the ultimatum game. *Games Econ Behav* 8:56–90
- Binmore K, Samuelson L, Vaughan R (1995b) Musical chairs: modeling noisy evolution. *Games Econ Behav* 11:1–35
- Binmore K, Samuelson L, Peyton Young H (2003) Equilibrium selection in bargaining models. *Games Econ Behav* 45:296–328
- Bishop DT, Cannings C (1978) A generalised war of attrition. *J Theor Biol* 70:85–124
- Bisin A, Verdier T (2001) The economics of cultural transmission and the dynamics of preferences. *J Econ Theory* 97:298–319
- Björnerstedt J, Weibull JW (1996) Nash equilibrium and evolution by imitation. In: Arrow KJ et al (eds) The rational foundations of economic behavior. St. Martin's Press, New York, pp 155–181
- Blume LE (1993) The statistical mechanics of strategic interaction. *Games Econ Behav* 5:387–424
- Blume LE (1995) The statistical mechanics of best response strategy revision. *Games Econ Behav* 11:111–145
- Blume LE (1997) Population games. In: Arthur WB, Durlauf SN, Lane DA (eds) The economy as an evolving complex system II. Addison-Wesley, Reading, pp 425–460
- Blume LE (2003) How noise matters. *Games Econ Behav* 44:251–271
- Bøg M (2006) Is segregation robust? Unpublished manuscript, Stockholm School of Economics
- Bomze IM (1990) Dynamical aspects of evolutionary stability. *Monatsh Math* 110:189–206
- Bomze IM (1991) Cross entropy minimization in uninhabitable states of complex populations. *J Math Biol* 30:73–87
- Börgers T, Sarin R (1997) Learning through reinforcement and the replicator dynamics. *J Econ Theory* 77:1–14
- Boylan RT (1995) Continuous approximation of dynamical systems with randomly matched individuals. *J Econ Theory* 66:615–625
- Brown GW, von Neumann J (1950) Solutions of games by differential equations. In: Kuhn HW, Tucker AW (eds) Contributions to the theory of games I. Annals of mathematics studies, vol 24. Princeton University Press, Princeton, pp 73–79
- Burke MA, Peyton Young H (2001) Competition and custom in economic contracts: a case study of Illinois agriculture. *Am Econ Rev* 91:559–573
- Cabrales A (1999) Adaptive dynamics and the implementation problem with complete information. *J Econ Theory* 86:159–184
- Cabrales A (2000) Stochastic replicator dynamics. *Int Econ Rev* 41:451–481
- Cabrales A, Ponti G (2000) Implementation, elimination of weakly dominated strategies and evolutionary dynamics. *Rev Econ Dyn* 3:247–282
- Crawford VP (1991) An “evolutionary” interpretation of Van Huyck, Battalio, and Beil's experimental results on coordination. *Games Econ Behav* 3:25–59
- Cressman R (1996) Evolutionary stability in the finitely repeated prisoner's dilemma game. *J Econ Theory* 68:234–248
- Cressman R (1997) Local stability of smooth selection dynamics for normal form games. *Math Soc Sci* 34:1–19

- Cressman R (2000) Subgame monotonicity in extensive form evolutionary games. *Games Econ Behav* 32:183–205
- Cressman R (2003) *Evolutionary dynamics and extensive form games*. MIT Press, Cambridge
- Cressman R, Schlag KH (1998) On the dynamic (in)stability of backwards induction. *J Econ Theory* 83:260–285
- Dafermos S, Sparrow FT (1969) The traffic assignment problem for a general network. *J Res Natl Bur Stand B* 73:91–118
- Dawid H, Bentley MacLeod W (2008) Hold-up and the evolution of investment and bargaining norms. *Games Econ Behav* 62:26–52. forthcoming
- Dawkins R (1976) *The selfish gene*. Oxford University Press, Oxford
- Dekel E, Scotchmer S (1992) On the evolution of optimizing behavior. *J Econ Theory* 57:392–407
- Demichelis S, Ritzberger K (2003) From evolutionary to strategic stability. *J Econ Theory* 113:51–75
- Dindoš M, Mezzetti C (2006) Better-reply dynamics and global convergence to Nash equilibrium in aggregative games. *Games Econ Behav* 54:261–292
- Dokumaci E, Sandholm WH (2007a) Schelling redux: an evolutionary model of residential segregation. Unpublished manuscript, University of Wisconsin
- Dokumaci E, Sandholm WH (2007b) Stochastic evolution with perturbed payoffs and rapid play. Unpublished manuscript, University of Wisconsin
- Droste E, Hommes C, Tuinstra J (2002) Endogenous fluctuations under evolutionary pressure in Cournot competition. *Games Econ Behav* 40:232–269
- Dugatkin LA, Reeve HK (eds) (1998) *Game theory and animal behavior*. Oxford University Press, Oxford
- Ellingsen T, Robles J (2002) Does evolution solve the hold-up problem? *Games Econ Behav* 39:28–53
- Ellison G (1993) Learning, local interaction, and coordination. *Econometrica* 61:1047–1071
- Ellison G (2000) Basins of attraction, long run equilibria, and the speed of step-by-step evolution. *Rev Econ Stud* 67:17–45
- Ely JC (2002) Local conventions. *Adv Econ Theory* 2:1(30)
- Ely JC, Sandholm WH (2005) Evolution in Bayesian games I: theory. *Games Econ Behav* 53:83–109
- Eshel I, Samuelson L, Shaked A (1998) Altruists, egoists, and hooligans in a local interaction model. *Am Econ Rev* 88:157–179
- Fischer S, Vöcking B (2006) On the evolution of selfish routing. Unpublished manuscript, RWTH Aachen
- Fisher RA (1930) *The genetical theory of natural selection*. Clarendon Press, Oxford
- Foster DP, Peyton Young H (1990) Stochastic evolutionary game dynamics. *Theor Popul Biol* 38:219–232. also in *Corrigendum* 51:77–78 (1997)
- Freidlin MI, Wentzell AD (1998) *Random perturbations of dynamical systems*, 2nd edn. Springer, New York
- Friedman D (1991) Evolutionary games in economics. *Econometrica* 59:637–666
- Friedman JW, Mezzetti C (2001) Learning in games by random sampling. *J Econ Theory* 98:55–84
- Friedman D, Yellin J (1997) Evolving landscapes for population games. Unpublished manuscript, UC Santa Cruz
- Fudenberg D, Harris C (1992) Evolutionary dynamics with aggregate shocks. *J Econ Theory* 57:420–441
- Fudenberg D, Imhof LA (2006) Imitation processes with small mutations. *J Econ Theory* 131:251–262
- Fudenberg D, Imhof LA (2008) Monotone imitation dynamics in large populations. *J Econ Theory* 140:229–245
- Fudenberg D, Levine DK (1998) *Theory of learning in games*. MIT Press, Cambridge
- Gaunersdorfer A, Hofbauer J (1995) Fictitious play, shapley polygons, and the replicator equation. *Games Econ Behav* 11:279–303
- Gilboa I, Matsui A (1991) Social stability and equilibrium. *Econometrica* 59:859–867
- Goyal S (2007) *Connections: an introduction to the economics of networks*. Princeton University Press, Princeton
- Goyal S, Janssen MCW (1997) Non-exclusive conventions and social coordination. *J Econ Theory* 77:34–57
- Hamilton WD (1967) Extraordinary sex ratios. *Science* 156:477–488
- Hammerstein P, Selten R (1994) Game theory and evolutionary biology, Chapter 28. In: Aumann RJ, Hart S (eds) *Handbook of game theory*, vol 2. Elsevier, Amsterdam, pp 929–993
- Harsanyi JC, Selten R (1988) *A general theory of equilibrium selection in games*. MIT Press, Cambridge
- Hart S (2002) Evolutionary dynamics and backward induction. *Games Econ Behav* 41:227–264
- Hart S, Mas-Colell A (2003) Uncoupled dynamics do not lead to Nash equilibrium. *Am Econ Rev* 93:1830–1836
- Hauert C (2007) *Virtual Labs in evolutionary game theory*. Software. <http://www.univie.ac.at/virtuallabs>. Accessed 31 Dec 2007
- Hauert C, De Monte S, Hofbauer J, Sigmund K (2002) Volunteering as red queen mechanism for cooperation in public goods games. *Science* 296:1129–1132
- Herz AVM (1994) Collective phenomena in spatially extended evolutionary games. *J Theor Biol* 169:65–87
- Hines WGS (1987) Evolutionary stable strategies: a review of basic theory. *Theor Popul Biol* 31:195–272
- Hofbauer J (1995a) Imitation dynamics for games. Unpublished manuscript, University of Vienna
- Hofbauer J (1995b) Stability for the best response dynamics. Unpublished manuscript, University of Vienna
- Hofbauer J (2000) From Nash and Brown to Maynard Smith: equilibria, dynamics and ESS. *Selection* 1:81–88
- Hofbauer J, Hopkins E (2005) Learning in perturbed asymmetric games. *Games Econ Behav* 52:133–152
- Hofbauer J, Sandholm WH (2002) On the global convergence of stochastic fictitious play. *Econometrica* 70:2265–2294
- Hofbauer J, Sandholm WH (2006a) Stable games. Unpublished manuscript, University of Vienna and University of Wisconsin
- Hofbauer J, Sandholm WH (2006b) Survival of dominated strategies under evolutionary dynamics. Unpublished manuscript, University of Vienna and University of Wisconsin
- Hofbauer J, Sandholm WH (2007) Evolution in games with randomly disturbed payoffs. *J Econ Theory* 132:47–69

- Hofbauer J, Sigmund K (1988) Theory of evolution and dynamical systems. Cambridge University Press, Cambridge
- Hofbauer J, Sigmund K (1998) Evolutionary games and population dynamics. Cambridge University Press, Cambridge
- Hofbauer J, Sigmund K (2003) Evolutionary game dynamics. *Bull Am Math Soc (New Ser)* 40:479–519
- Hofbauer J, Swinkels JM (1996) A universal Shapley example. Unpublished manuscript, University of Vienna and Northwestern University
- Hofbauer J, Weibull JW (1996) Evolutionary selection against dominated strategies. *J Econ Theory* 71:558–573
- Hofbauer J, Schuster P, Sigmund K (1979) A note on evolutionarily stable strategies and game dynamics. *J Theor Biol* 81:609–612
- Hofbauer J, Oechssler J, Riedel F (2005) Brown-von Neumann-Nash dynamics: the continuous strategy case. Unpublished manuscript, University of Vienna
- Hopkins E (1999) A note on best response dynamics. *Games Econ Behav* 29:138–150
- Hopkins E, Seymour RM (2002) The stability of price dispersion under seller and consumer learning. *Int Econ Rev* 43:1157–1190
- Imhof LA (2005) The long-run behavior of the stochastic replicator dynamics. *Ann Appl Probab* 15:1019–1045
- Jackson MO (2017) Social and economic networks. Princeton University Press, Princeton
- Jacobsen HJ, Jensen M, Sloth B (2001) Evolutionary learning in signalling games. *Games Econ Behav* 34:34–63
- Jordan JS (1993) Three problems in learning mixed-strategy Nash equilibria. *Games Econ Behav* 5:368–386
- Josephson J (2008) Stochastic better reply dynamics in finite games. *Econ Theory* 35:381–389
- Josephson J, Matros A (2004) Stochastic imitation in finite games. *Games Econ Behav* 49:244–259
- Kandori M, Rob R (1995) Evolution of equilibria in the long run: a general theory and applications. *J Econ Theory* 65:383–414
- Kandori M, Rob R (1998) Bandwagon effects and long run technology choice. *Games Econ Behav* 22:84–120
- Kandori M, Mailath GJ, Rob R (1993) Learning, mutation, and long run equilibria in games. *Econometrica* 61:29–56
- Kim Y-G, Sobel J (1995) An evolutionary approach to pre-play communication. *Econometrica* 63:1181–1193
- Kimura M (1958) On the change of population fitness by natural selection. *Heredity* 12:145–167
- Kosfeld M (2002) Stochastic strategy adjustment in coordination games. *Econ Theory* 20:321–339
- Kukushkin NS (2004) Best response dynamics in finite games with additive aggregation. *Games Econ Behav* 48:94–110
- Kuran T, Sandholm WH (2008) Cultural integration and its discontents. *Rev Econ Stud* 75:201–228
- Kurtz TG (1970) Solutions of ordinary differential equations as limits of pure jump Markov processes. *J Appl Probab* 7:49–58
- Kuzmics C (2004) Stochastic evolutionary stability in extensive form games of perfect information. *Games Econ Behav* 48:321–336
- Lahkar R (2007) The dynamic instability of dispersed price equilibria. Unpublished manuscript, University College London
- Lahkar R, Sandholm WH (2017) The projection dynamic and the geometry of population games. *Games Econ Behav*
- Losert V, Akin E (1983) Dynamics of games and genes: discrete versus continuous time. *J Math Biol* 17:241–251
- Lotka AJ (1920) Undamped oscillation derived from the law of mass action. *J Am Chem Soc* 42:1595–1598
- Mailath GJ (1992) Introduction: symposium on evolutionary game theory. *J Econ Theory* 57:259–277
- Maruta T (1997) On the relationship between risk-dominance and stochastic stability. *Games Econ Behav* 19:221–234
- Maruta T (2002) Binary games with state dependent stochastic choice. *J Econ Theory* 103:351–376
- Mathevet L (2007) Supermodular Bayesian implementation: learning and incentive design. Unpublished manuscript, Caltech
- Maynard Smith J (1972) Game theory and the evolution of fighting. In: Maynard Smith J (ed) *On evolution*. Edinburgh University Press, Edinburgh, pp 8–28
- Maynard Smith J (1974) The theory of games and the evolution of animal conflicts. *J Theor Biol* 47:209–221
- Maynard Smith J (1982) *Evolution and the theory of games*. Cambridge University Press, Cambridge
- Maynard Smith J, Price GR (1973) The logic of animal conflict. *Nature* 246:15–18
- Miękisz J (2004) Statistical mechanics of spatial evolutionary games. *J Phys A* 37:9891–9906
- Möbius MM (2000) The formation of ghettos as a local interaction phenomenon. Unpublished manuscript, MIT
- Monderer D, Shapley LS (1996) Potential games. *Games Econ Behav* 14:124–143
- Moran PAP (1962) *The statistical processes of evolutionary theory*. Clarendon Press, Oxford
- Myatt DP, Wallace CC (2003) A multinomial probit model of stochastic evolution. *J Econ Theory* 113:286–301
- Myatt DP, Wallace CC (2007) An evolutionary justification for thresholds in collective-action problems. Unpublished manuscript, Oxford University
- Myatt DP, Wallace CC (2008a) An evolutionary analysis of the volunteer's dilemma. *Games Econ Behav* 62:67–76
- Myatt DP, Wallace CC (2008b) When does one bad apple spoil the barrel? An evolutionary analysis of collective action. *Rev Econ Stud* 75:499–527
- Nachbar JH (1990) “Evolutionary” selection dynamics in games: convergence and limit properties. *Int J Game Theory* 19:59–89
- Nagurney A, Zhang D (1997) Projected dynamical systems in the formulation, stability analysis and computation of fixed demand traffic network equilibria. *Transp Sci* 31:147–158
- Nash JF (1951) Non-cooperative games. *Ann Math* 54:287–295

- Nöldeke G, Samuelson L (1993) An evolutionary analysis of backward and forward induction. *Games Econ Behav* 5:425–454
- Nowak MA (2006) *Evolutionary dynamics: exploring the equations of life*. Belknap/Harvard, Cambridge
- Nowak MA, May RM (1992) Evolutionary games and spatial chaos. *Nature* 359:826–829
- Nowak MA, May RM (1993) The spatial dilemmas of evolution. *Int J Bifurcat Chaos* 3:35–78
- Nowak MA, Bonhoeffer S, May RM (1994a) More spatial games. *Int J Bifurcat Chaos* 4:33–56
- Nowak MA, Bonhoeffer S, May RM (1994b) Spatial games and the maintenance of cooperation. *Proc Natl Acad Sci U S A* 91:4877–4881
- Nowak MA, Sasaki A, Taylor C, Fudenberg D (2004) Emergence of cooperation and evolutionary stability in finite populations. *Nature* 428:646–650
- Oechssler J, Riedel F (2001) Evolutionary dynamics on infinite strategy spaces. *Econ Theory* 17:141–162
- Oechssler J, Riedel F (2002) On the dynamic foundation of evolutionary stability in continuous models. *J Econ Theory* 107:141–162
- Peyton Young H (1993a) The evolution of conventions. *Econometrica* 61:57–84
- Peyton Young H (1993b) An evolutionary model of bargaining. *J Econ Theory* 59:145–168
- Peyton Young H (1998a) Conventional contracts. *Rev Econ Stud* 65:773–792
- Peyton Young H (1998b) *Individual strategy and social structure*. Princeton University Press, Princeton
- Peyton Young H (2001) The dynamics of conformity. In: Durlauf SN, Peyton Young H (eds) *Social dynamics*. Brookings Institution Press/MIT Press, Washington, DC/Cambridge, pp 133–153
- Rhode P, Stegeman M (1996) A comment on “learning, mutation, and long run equilibria in games”. *Econometrica* 64:443–449
- Ritzberger K, Weibull JW (1995) Evolutionary selection in normal form games. *Econometrica* 63:1371–1399
- Robles J (1998) Evolution with changing mutation rates. *J Econ Theory* 79:207–223
- Robles J (2008) Evolution, bargaining and time preferences. *Econ Theory* 35:19–36
- Robson A, Vega-Redondo F (1996) Efficient equilibrium selection in evolutionary games with random matching. *J Econ Theory* 70:65–92
- Rosenthal RW (1973) A class of games possessing pure strategy Nash equilibria. *Int J Game Theory* 2:65–67
- Samuelson L (1988) Evolutionary foundations of solution concepts for finite, two-player, normal-form games. In: Vardi MY (ed) *Proceedings of the second conference on theoretical aspects of reasoning about knowledge* (Pacific Grove, CA, 1988). Morgan Kaufmann Publishers, Los Altos, pp 211–225
- Samuelson L (1994) Stochastic stability in games with alternative best replies. *J Econ Theory* 64:35–65
- Samuelson L (1997) *Evolutionary games and equilibrium selection*. MIT Press, Cambridge
- Samuelson L, Zhang J (1992) Evolutionary stability in asymmetric games. *J Econ Theory* 57:363–391
- Sandholm WH (1998) Simple and clever decision rules in a model of evolution. *Econ Lett* 61:165–170
- Sandholm WH (2001a) Almost global convergence to p -dominant equilibrium. *Int J Game Theory* 30:107–116
- Sandholm WH (2001b) Potential games with continuous player sets. *J Econ Theory* 97:81–108
- Sandholm WH (2002) Evolutionary implementation and congestion pricing. *Rev Econ Stud* 69:81–108
- Sandholm WH (2003) Evolution and equilibrium under inexact information. *Games Econ Behav* 44:343–378
- Sandholm WH (2005a) Excess payoff dynamics and other well-behaved evolutionary dynamics. *J Econ Theory* 124:149–170
- Sandholm WH (2005b) Negative externalities and evolutionary implementation. *Rev Econ Stud* 72:885–915
- Sandholm WH (2006) Pairwise comparison dynamics. Unpublished manuscript, University of Wisconsin
- Sandholm WH (2007a) Evolution in Bayesian games II: stability of purified equilibria. *J Econ Theory* 136:641–667
- Sandholm WH (2007b) Pigouvian pricing and stochastic evolutionary implementation. *J Econ Theory* 132:367–382
- Sandholm WH (2007c) Large population potential games. Unpublished manuscript, University of Wisconsin
- Sandholm WH (2007d) Simple formulas for stationary distributions and stochastically stable states. *Games Econ Behav* 59:154–162
- Sandholm WH (2017) *Population games and evolutionary dynamics*. MIT Press, Cambridge
- Sandholm WH, Dokumaci E (2007) *Dynamo: phase diagrams for evolutionary dynamics*. Software. <http://www.ssc.wisc.edu/~whs/dynamo>
- Sandholm WH, Pauzner A (1998) Evolution, population growth, and history dependence. *Games Econ Behav* 22:84–120
- Sandholm WH, Dokumaci E, Lahkar R (2017) The projection dynamic and the replicator dynamic. *Games Econ Behav*
- Sato Y, Akiyama E, Doyne Farmer J (2002) Chaos in learning a simple two-person game. *Proc Natl Acad Sci U S A* 99:4748–4751
- Schlag KH (1998) Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *J Econ Theory* 78:130–156
- Schuster P, Sigmund K (1983) Replicator dynamics. *J Theor Biol* 100:533–538
- Schuster P, Sigmund K, Hofbauer J, Wolff R (1981) Selfregulation of behaviour in animal societies I: symmetric contests. *Biol Cybern* 40:1–8
- Selten R (1991) Evolution, learning, and economic behavior. *Games Econ Behav* 3:3–24
- Shahshahani S (1979) A new mathematical framework for the study of linkage and selection. *Mem Am Math Soc* 211:34
- Shapley LS (1964) Some topics in two person games. In: Dresher M, Shapley LS, Tucker AW (eds) *Advances in game theory*. Annals of mathematics studies, vol 52. Princeton University Press, Princeton, pp 1–28

- Skyrms B (1990) The dynamics of rational deliberation. Harvard University Press, Cambridge
- Skyrms B (1992) Chaos in game dynamics. *J Log Lang Inf* 1:111–130
- Smith MJ (1984) The stability of a dynamic model of traffic assignment -an application of a method of Lyapunov. *Transp Sci* 18:245–252
- Smith HL (1995) Monotone dynamical systems: an introduction to the theory of competitive and cooperative systems. American Mathematical Society, Providence
- Stegeman M, Rhode P (2004) Stochastic Darwinian equilibria in small and large populations. *Games Econ Behav* 49:171–214
- Swinkels JM (1992) Evolutionary stability with equilibrium entrants. *J Econ Theory* 57:306–332
- Swinkels JM (1993) Adjustment dynamics and rational play in games. *Games Econ Behav* 5:455–484
- Szabó G, Fáth G (2007) Evolutionary games on graphs. *Phys Rep* 446:97–216
- Szabó G, Hauert C (2002) Phase transitions and volunteering in spatial public goods games. *Phys Rev Lett* 89:11801(4)
- Tainaka K-I (2001) Physics and ecology of rock-paper-scissors game. In: Marsland TA, Frank I (eds) Computers and games, second international conference (Hamamatsu 2000). Lecture notes in computer science, vol 2063. Springer, Berlin, pp 384–395
- Tanabe Y (2006) The propagation of chaos for interacting individuals in a large population. *Math Soc Sci* 51:425–152
- Taylor PD, Jonker L (1978) Evolutionarily stable strategies and game dynamics. *Math Biosci* 40:145–156
- Thomas B (1985) On evolutionarily stable sets. *J Math Biol* 22:105–115
- Topkis D (1979) Equilibrium points in nonzero-sum n -person submodular games. *SIAM J Control Optim* 17:773–787
- Tröger T (2002) Why sunk costs matter for bargaining outcomes: an evolutionary approach. *J Econ Theory* 102:28–53
- Ui T (1998) Robustness of stochastic stability. Unpublished manuscript, Bank of Japan
- van Damme E, Weibull JW (2002) Evolution in games with endogenous mistake probabilities. *J Econ Theory* 106:296–315
- Vega-Redondo F (1996) Evolution, games, and economic behaviour. Oxford University Press, Oxford
- Vega-Redondo F (1997) The evolution of Walrasian behavior. *Econometrica* 65:375–384
- Vega-Redondo F (2007) Complex social networks. Cambridge University Press, Cambridge
- Volterra V (1931) Lecons sur la Theorie Mathematique de la Lutte pour la Vie. Gauthier-Villars, Paris
- von Neumann J, Morgenstern O (1944) Theory of games and economic behavior. Prentice-Hall, Princeton
- Weibull JW (1995) Evolutionary game theory. MIT Press, Cambridge
- Weibull JW (1996) The mass action interpretation. Excerpt from “The work of John Nash in game theory: nobel seminar, December 8, 1994”. *J Econ Theory* 69:165–171
- Weissing FJ (1991) Evolutionary stability and dynamic stability in a class of evolutionary normal form games. In: Selten R (ed) Game equilibrium models I. Springer, Berlin, pp 29–97
- Zeeman EC (1980) Population dynamics from game theory. In: Nitecki Z, Robinson C (eds) Global theory of dynamical systems (Evanston, 1979). Lecture notes in mathematics, vol 819. Springer, Berlin, pp 472–497
- Zhang J (2004a) A dynamic model of residential segregation. *J Math Sociol* 28:147–170
- Zhang J (2004b) Residential segregation in an all-integrationist world. *J Econ Behav Organ* 24:533–550

Networks and Stability

Frank H. Page Jr.¹ and Myrna Wooders²

¹Department of Economics, Indiana University,
Bloomington, IN, USA

²Department of Economics, Vanderbilt
University, Nashville, TN, USA

Article Outline

Glossary

Definition of the Subject

Introduction

The Primitives

Abstract Games of Network Formation and
Stability

Strong Stability, Pairwise Stability, Nash Stability,
and Farsighted Consistency

Singleton Basins of Attraction

Future Directions

Bibliography

Glossary

Abstract game of network formation with respect to irreflexive dominance An abstract game of network formation with respect to irreflexive dominance consists of a feasible set of networks G equipped with a irreflexive dominance relation $>$. A dominance relation on G is a binary relation on G such that for all G and G' in G , $G' > G$ (read G' dominates G) is either true or false. The dominance relation is irreflexive if $G > G$ is always false.

Abstract game of network formation with respect to path dominance An abstract game of network formation with respect to path dominance consists of a feasible set of networks G equipped with a path dominance relation $\geq_p$ induced by an irreflexive

dominance relation $>$ on G . Given networks G and G' in G , $G' \geq_p G$ (read G' path dominates G) if either $G' = G$ or there is a finite sequence of networks in G beginning with G and ending with G' such that each network along the sequence dominates its predecessor.

Heterogeneous networks A heterogeneous network consists of a finite set of nodes together with a finite set of mathematical objects called labeled links or labeled arcs, each identifying a particular type of connection between a pair of nodes. Given finite node set N with typical element i and given finite label set A with typical element a , a heterogeneous linking network G is a finite collection of ordered pairs of the form $(a, \{i, i'\})$ called labeled links. Labeled link $(a, \{i, i'\}) \in G$ indicates that nodes i and i' are connected in network G via a type a link. A heterogeneous directed network G is a finite collection of ordered pairs of the form $(a, (i, i'))$ called labeled arcs. Labeled arc $(a, (i, i')) \in G$ indicates that nodes i and i' are connected in network G via a type a arc running from i to i' . In a heterogeneous network (whether it be a linking network or a directed network) connections can differ and are distinguished by type.

Homogeneous networks A homogeneous network consists of a finite set of nodes together with a finite set of mathematical objects called links or arcs, each identifying a connection between a pair of nodes. Given finite node set N with typical element i , a homogeneous linking network G is a finite collection of sets of the form $\{i, i'\}$ called links. Link $\{i, i'\} \in G$ indicates that nodes i and i' are connected in network G . A homogeneous directed network G is a finite collection of ordered pairs (i, i') called arcs. Arc $(i, i') \in G$ indicates that nodes i and i' are connected in network G via a connection running from i to i' . In a homogeneous network (whether it be a linking network or a directed network) all connections are of the same type.

Definition of the Subject

Our subject is networks, and in particular, stable networks and the game theoretic underpinnings of stable networks.

Networks are pervasive. We routinely communicate over the internet, advance our careers by networking, travel to conferences over the transportation network and pay for the trip using the banking network. Doing this utilizes networks in our brain. The list could go on. While network models have had a long history in sociology, the natural sciences, and engineering (e. g., in modeling social organizations, brain architecture, and electrical circuits), the rise of the network paradigm in economics is relatively recent. Economists are now beginning to think of political and economic interactions as network phenomena and to model everything from terrorist activities to asset market micro structures as *games of network formation*. This trend in economics, which began with the seminal paper by Myerson (1977) on graphs and cooperation and accelerated with the publication of the papers by Jackson and Wolinsky (1996) and Dutta and Mutuswami (1997) on stable and efficient networks, is likely to continue with the development of new algorithms, the expansion of computational capacity and the broad application of network theories to economic, political, and social phenomena.

What economists bring to the study of networks that is new is game theory. For the most part sociologists, natural scientists and engineers have used networks descriptively and have focused on the design of networks from the perspective of a single designer or on the random evolution of networks from the perspective of nature. This singularity of perspective is a consequence of the nonstrategic nature of the phenomena being explained or the problem being solved (e.g., the spread of a disease through a given population, the transmission of electrical impulses in the brain, or the optimal design of an integrated circuit). In economics the perspective is often times strategic. In particular, in many economic situations, several individuals, guided by their own self interest, behave strategically in putting into place pieces of the network of economic,

political, or social interactions under their control and in so doing generate payoffs and externalities that determine the network of economic interactions that eventually emerges in equilibrium. Thus in economics, pieces of the network are the strategies and the network that ultimately prevails is the result of strategic competition rather than the design of a single individual or nature. Because large computer networks such as the internet are built, operated, and used by a many diverse individuals with competing interests, computer scientists are also beginning to use game theoretic models to analyze and understand the optimal design of secure computer networks (see for example, Roughgarden 2005; Tardos and Wexler 2007).

Conversely, what networks bring to the study of economics is a new way of modeling the structure of economic interactions and externalities that makes possible a game-theoretic analysis of how these structures influence individual payoffs and the economic equilibrium that emerges from competition.

Introduction

Our main objective is to present a unified, game-theoretic development of the main concepts of stability that have appeared in the recent economics literature on strategic network formation, specifically, the notions of strong stability (Jackson and van den Nouweland 2005), pairwise stability (Jackson and Wolinsky 1996), Nash stability, and farsighted consistency (Chwe 1994). In order to accomplish this we follow the approach introduced in Page and Wooders (2005, 2008). The key ingredient in this approach is an abstract game model of endogenous network formation (i.e., abstract game in the sense of von Neumann-Morgenstern (1944)). The model is built on four primitives: A feasible set of networks, player preferences over networks, the rules of network formation, and an irreflexive dominance relation over networks. In the remainder of this introduction we provide an overview of the four primitives of our model, a summary of results discussed, and note some

important areas of research that are beyond the scope of this entry.

Feasible Sets The feasible set may consist of networks as simple as homogenous linking networks or as complex as heterogeneous directed networks. All networks consist of a finite set of nodes (representing, for example, economic agents or players) together with a finite set of mathematical objects called links, labeled links, arcs, or labeled arcs describing the connections between nodes. Here we will focus on homogeneous linking networks, as does most of literature (see, for example, Myerson (1977), Jackson and Wolinsky (1996), and Jackson and van den Nouweland (2005)), except in our discussion of Nash stability where we will consider homogeneous directed networks (as in Bala and Goyal 2000). What distinguishes homogeneous networks (linking or directed) from heterogeneous networks (linking or directed) is that in a homogeneous network all connections between nodes are of the same type whether represented by a link as in a linking network or by an arc as in a directed network. Thus, in a homogeneous linking network all links are of the same type and in a homogeneous directed network all arcs are of the same type. While homogeneous networks are quite restrictive, they have been very important in developing our understanding of social and economic networks and have proved very useful in many economic applications (see, for example, Belleflamme and Bloch (2004), Bramoulle and Kranton (2007a), Calvo-Armengol (2004), and Furusawa and Konishi (2007)). Page and Wooders (2005, 2008; Page and Kamat 2005; Page et al. 2005) extend the existing literature on economic and social networks by introducing the notion of heterogeneous directed networks. These types of networks potentially have a rich set applications (in the natural sciences, engineering, sociology, politics, as well as economics) because connections or interactions between nodes can be distinguished by direction or intent as well as by type, intensity, or purpose.

Players' Preferences We will assume throughout that each player's preferences are given by an

irreflexive binary relation defined on the feasible set of networks. Thus, we will assume that players have strong (or strict) preferences over networks. Under strong preferences, if a player prefers one network to another, then the player's preference is strict. However, we will comment where appropriate on weak preferences. Under weak preferences, if a player prefers one network to another, then the player's preference is either strict or indifferent.

Rules of Network Formation We will focus here on three different sets of rules: Jackson-Wolinsky rules (Jackson and Wolinsky 1996), Jackson-van den Nouweland rules (Jackson and van den Nouweland 2005), and Bala-Goyal rules (Bala and Goyal 2000). In particular, in our discussions of pairwise stable homogeneous linking networks we will assume that the rules of network formation are the Jackson-Wolinsky rules. Under the Jackson-Wolinsky rules the addition of a link is bilateral (i.e., the two players that would be involved in the link must agree to adding the link), the subtraction of a link is unilateral (i.e., at least one player involved in the link must agree to subtract or delete the link), and network changes take place one link at a time (i.e., only one link can be added or subtracted at a time). In our discussion of strongly stable homogeneous linking networks, we will assume that the rules of network formation are the Jackson-van den Nouweland rules. Under the Jackson-van den Nouweland rules link addition is bilateral, link subtraction is unilateral, and in any one play of the game several links can be added and/or subtracted. Thus the Jackson-van den Nouweland rules are the Jackson-Wolinsky rules without the one-link-at-a-time restriction. Finally, in our discussion of Nash homogeneous directed networks we will assume that the rules of network formation are the Bala-Goyal rules. Under the Bala-Goyal rules an arc may be added or subtracted unilaterally by the initiating player involved in the arc and in any one play of the game only network changes brought about by an individual player are allowed. Note that all three of these sets of rules can be described as being uniform across networks. Under uniform rules the rules for changing a network are the same no matter which status quo

network is being changed. Page et al. (2005) allow nonuniform rules and introduce a network representation of nonuniform rules.

Dominance Relations Given players' preferences and the rules of network formation we will define a dominance relation over the feasible set of networks that incorporates both players preferences and the rules. Here we will focus on dominance relations that are either direct or indirect. Under direct dominance players are concerned with immediate consequences of their network formation strategies whereas under indirect dominance players are farsighted and consider the eventual consequences of their strategies.

General Results A specification of the primitives induces two types of abstract games over homogeneous networks: (i) a network formation game with respect to the irreflexive dominance relation induced by preferences and rules, and (ii) a network formation game with respect to *path* dominance induced by this irreflexive dominance relation. We will begin by considering the game with respect to irreflexive dominance and present results on the existence of quasi-stable and stable networks. These results provide a network rendition of classical results from graph theory on the existence of quasi-stable sets and stable sets due to Chvatal and Lovasz (1972), Berge (2001), and Richardson (1953). We will also present a result on the existence and nonemptiness of the set of farsightedly consistent networks. This result is a network rendition of a result due to Chwe (1994) for abstract games.

Next we will consider the game over homogeneous networks with respect to path dominance, and we will conclude that the following results hold:

1. Given preferences and the rules governing network formation, the set of homogeneous networks (linking or directed) contains a unique, finite, disjoint collection of nonempty subsets each constituting a *strategic basin of attraction*. These basins of attraction are the absorbing sets of the competitive process of network formation modeled via the game.

2. A stable set of homogeneous networks (in the sense of von Neumann-Morgenstern) with respect to path dominance consists of one network from each basin of attraction.
3. The path dominance core, defined as the set networks having the property that no network in the set is path dominated by any other homogeneous network, consists of one network from each basin of attraction containing a *single* network. Note that the path dominance core is contained in each stable set and is nonempty if and only if there is a basin of attraction containing a single network. As a corollary, we conclude that any homogeneous network contained in the path dominance core is constrained Pareto efficient.
4. From the results above it follows that if the dominance relation is transitive and irreflexive, then the path dominance core is nonempty.

These results are a special cases of results due to Page and Wooders (2008).

Specific Results for Pairwise Stability, Strong Stability, Nash Stability, and Farsighted Consistency

What are the connections between our notions of stability for homogeneous networks (basins of attraction, path dominance stable sets, and path dominance core) and the notions of strong stability (Dutta and Mutuswami 1997; Jackson and van den Nouweland 2005), pairwise stability (Jackson and Wolinsky 1996), Nash stability (Bala and Goyal 2000), and farsighted consistency (Chwe 1994; Page et al. 2005)? From the general results in (Page and Wooders 2005; Page and Wooders 2008) for heterogeneous directed networks, we will conclude for the case of homogeneous networks (linking or directed) that, depending on how we specialize the primitives of the model, the path dominance core is equal to the set of strongly stable networks, the set of pairwise stable networks, or the set of Nash networks. In particular, we will conclude that:

1. If path dominance is induced by a direct dominance relation, then in the set of homogeneous linking networks the path dominance

- core is equal to the set of strongly stable networks.
2. If, in addition, the rules of network formation are the Jackson-Wolinsky rules, then in the set of homogeneous linking networks the path dominance core is equal to the set of pairwise stable networks.
 3. If path dominance is induced by a direct dominance relation and if the rules of network formation are the Bala-Goyal rules, then in the set of homogeneous directed networks the path dominance core is equal to the set of Nash networks.

We can then conclude from (3) above that the existence of at least one basin of attraction containing a single network is, depending on how we specialize primitives, both necessary and sufficient for either (i) the existence of a strongly stable network, or (ii) a pairwise stable network, or (iii) a Nash network.

For path dominance induced by an indirect dominance relation, we can conclude from our prior results that for the case of homogeneous linking networks with Jackson-Wolinsky or Jackson-van den Nouweland rules or for the case of homogeneous directed networks with Bala-Goyal rules, each strategic basin of attraction has a nonempty intersection with the largest farsightedly consistent set of networks. This result, together with (2) above, implies that there always exists a path dominance stable set of homogeneous networks contained in the largest farsightedly consistent set. Thus, the path dominance core is contained in the largest consistent set. In light of our results on the path dominance core and stability (both strong and pairwise), we conclude that if path dominance is induced by an indirect dominance relation, then any homogeneous network contained in the path dominance core (i.e., the farsighted core) is not only farsightedly consistent but also strongly stable, as well as pairwise stable. Other papers using indirect dominance (or variations thereof) and farsighted consistency in games (not necessarily network formation games) include Li (1992, 1993), Xue (1998, 2000), Luo (2001), Mariotti and Xue (2002), Diamantoudi and Xue (2003), and Mauleon and

Vannetelbosch (2004), Bhattacharya (2005), Herrings et al. (2006).

We remark that solution concepts defined using dominance relations have a long and distinguished history in the literature of game theory. First, consider the von Neuman-Morgenstern stable set (see Richardson 1953; von Neumann and Morgenstern 1944). The vN-M stable set is defined with respect to a dominance relation on a set of outcomes and consists of those outcomes that are externally and internally stable with respect to the given dominance relation. Similarly, Gillies (1959) defines the core based on a given dominance relation. These solution concepts, with a few exceptions, have typically been applied to models of economies or cooperative games where the notion of dominance is based on what a coalition can achieve using only the resources owned by its members (cf., Aumann 1964) or a given set of utility vectors for each possible coalition (cf., Scarf 1967). Particularly notable exceptions are Schwartz (1974), Kalai et al. (1976), Kalai and Schmeidler (1977), Shenoy (1980), Inarra et al. (2005), and van Deemen (1991). Their motivations are in part similar to ours in that they take as given a set of possible choices for players (here consisting of set of networks) and a dominance relation and, based on these, describe a set of possible or likely outcomes called, by Kalai and Schmeidler, the admissible set. While their examples treat direct dominance, their general results have wider applications.

Because our objective here is to provide a unified game theoretic treatment of the main stability notions for network formation games, many topics related to strategic networks are not covered here. For example, we do not discuss the conflict between stability and efficiency which is the main focus of the important papers by Dutta and Mutuswami (1997) and Currarini and Morelli (2000), and Mutuswami and Winter (2002). Nor do we treat the topic of network formation and cooperative games, the topic of the seminal paper by Myerson (1977) and the excellent book by Slikker and van den Nouweland (2001) among many other contributions, or the topic of network formation and evolution treated in Hojman and Szeidl (2006). Our game theoretic approach

provides a snapshot of all possible network formation paths under dynamics which respect preferences and the rules of network formation. Our approach, however, is not explicitly dynamic. For network dynamics, we can only suggest to the reader the elegant papers by Skyrms and Pemantle (2000), Watts (2001), Jackson and Watts (2002), Konishi and Ray (2003), and Dutta et al. (2005). We do not touch on the topic of learning in networks, a topic which has been the focus of much work by Goyal (2005, 2007), nor do we discuss random networks, introduced in economics by Kirman (1983). For random networks we refer the reader to the recent book by Vega-Redondo (2007) and the references contained therein. Finally, we do not discuss the statistical mechanics of network formation. For this topic, we recommend to the reader the excellent papers by Blume (1993) and Durlauf (1997).

Because our focus is on foundational issues in strategic network formation, and in particular stability, we do not discuss any of the plethora of economic applications that can be found in the exploding literature on social and economic networks. For now, we can only offer the reader the following modest and incomplete list of topics and papers:

Development and insurance	Bloch et al. (2008), Bramoulle and Kranton (2007b)
Employment and labor markets	Rees (1966), Granovetter (1973), Boorman (1975), Montgomery (1991), Topa (2001), Calvo-Armengol (2004), Calvo-Armengol and Jackson (2004), Calvo-Armengol and Jackson (2007)
Industrial organization and R&D	Demange and Henriët (1991), Bloch (1995), Kranton and Minehart (2000), Goyal and Moraga-Gonzalez (2001), Goyal and Joshi (2003), Belleflamme and Bloch (2004), Bloch (2005), and Deroian and Gannon (2005), Mauleon et al. (2008), Wang and Watts (2006)
International trade	Casella and Rauch (2002, 2003) Zissimos (Arnold and Wooders 2006), Goyal and

(continued)

	Joshi (2006), and Furusawa and Konishi (2007);
Market microstructure	Tesfatsion (1997, 1998), Kirman et al. (2000), Kranton and Minehart (2001), Corominas-Bosch (2004), and Even-Dar et al. (2007)
Public goods	Bramoulle and Kranton (2007a)
Organizations, coordination, and communication	Chwe (2000), Currarini (2007), Demange (2004)
Marketing and advertising	Marketing and Galeotti and Moraga-Gonzalez (2007)

The Primitives

Our abstract game of network formation rests on four primitives: The feasible set of networks, players’ preferences, the rules of network formation, and a dominance relation over feasible networks. In this section, we discuss in detail these four primitives and in the next section, using these primitives we construct our abstract games of network formation.

Feasible Networks

Types of Networks

Surprisingly, there is no agreed-upon definition of a network but rather several definitions depending on the application. But what all definitions have in common is a nonempty set of nodes and a precise mathematical description of how nodes are connected. What differentiates these various definitions then are the details of how nodes are connected. We begin with the most elementary notion of a network, the homogeneous linking network, and proceed to a more complex notion, the heterogeneous directed network introduced in (Page and Wooders 2005). Let N be a finite set of nodes, with typical element denoted by i , and let A be a finite set of link types or arc types, with typical element denoted by a . If the network is directed (to be defined below), we refer to the elements of A as arc types, otherwise we will refer to the elements of A as link types. For any set E , we denote by $P(E)$ the collection of all

subsets of E . Finally, for any set E , we denote by $|E|$ the cardinality of E (note that $|\emptyset| = 0$).

Linking Networks Definition 1 (Homogeneous Linking Networks, Myerson (1977), Jackson-Wolinsky (1996)) Let $P_2(N)$ denote the set of all subsets of N of size 2. A linking network, G , is a subset (possibly empty) of $P_2(N)$ and for any $G \subseteq P_2(N)$, each subset $\{i, i'\} \in G$ is called a link in G . The collection of all homogeneous linking networks is denoted by $P(P_2(N))$.

Thus, $P_2(N)$ is the set of all possible links and a homogeneous linking network G is simply a subset of all possible links. For example, if $N = \{i_1, i_2, i_3\}$, then

$$P_2(N) = \{\{i_1, i_2\}, \{i_2, i_3\}, \{i_1, i_3\}\},$$

and the subset


$$G_1 = \{\{i_1, i_2\}, \{i_1, i_3\}\}$$

of $P_2(N)$ is a homogeneous linking network. Figure 1 depicts homogeneous linking network G_1 . Here, the link $\{i_1, i_3\} \in G_1$ denotes that nodes i_1 and i_3 are connected or linked. Note that all links are the same (i.e., links are homogeneous) and links have no orientation or direction. Also, note that in a homogeneous linking network, loops are not allowed by definition (a loop being a link between a node and itself). Finally, note that in a linking network multiple links between any pair of nodes are not allowed. However, because links are homogeneous, multiple links are unnecessary.


The following extended definition of a linking network allows for heterogeneous links. This heterogeneity is represented by a labeling of links using elements of the set A of link types.

Definition 2 (Heterogeneous Linking Networks) A heterogeneous linking network, G , is a subset of $A \times P_2(N)$. Given any $G \subseteq A \times P_2(N)$, each ordered pair $(a, \{i, i'\}) \in G$ consisting of a link type and a link is called a labeled link in G . The collection of all heterogeneous linking networks is denoted by $P(A \times P_2(N))$.

Thus, $A \times P_2(N)$ is the set of all possible labeled links and a heterogeneous linking network G is simply a subset of all possible labeled links.



Networks and Stability, Fig. 1 Homogeneous linking network G_1



Networks and Stability, Fig. 2 Heterogeneous linking network G_2

For example, given $N = \{i_1, i_2, i_3\}$ and $A = \{a_1, a_2, a_3\}$, the subset

$$G_2 = \{(a_1, \{i_1, i_2\}), (a_1, \{i_1, i_3\}), (a_3, \{i_2, i_3\})\}$$

of $A \times P_2(N)$ is a heterogeneous linking network. Figure 2 depicts heterogeneous linking network G_2 .

Here, the labeled link $(a_1, \{i_1, i_3\}) \in G_2$ denotes that nodes i_1 and i_3 are linked by a type a_1 link. First, note that in network G_2 in addition to being linked by an a_1 link, nodes i_1 and i_3 are also linked by an a_3 link. Thus, in a heterogeneous linking network, links are not identical and multiple, distinct links between any given pair of nodes are possible. Second, note that in network G_2 , nodes i_1 and i_2 – like nodes i_1 and i_3 – are linked by an a_1 link. Thus, in a heterogeneous linking network, link types can be used multiple times for different pairs of nodes. Finally, note that

in a heterogeneous linking network, links are still without orientation or direction and loops are not possible.

Directed Networks The link orientation problem as well as the problem of loops is resolved by moving to directed networks. As is the case with linking networks, there are two categories of directed networks: Homogeneous directed networks and heterogeneous directed networks.

Definition 3 (Homogeneous Directed Networks) A homogeneous directed network, G , is a subset of $N \times N$. Given any $G \subseteq N \times N$, each ordered pair $(i, i') \in G$ consisting of a beginning node i and an ending node i' is called an arc in G . The collection of all directed networks is denoted by $P(N \times N)$.

Thus, $N \times N$ is the set of all possible arcs and a homogeneous directed network G is simply a subset of all possible arcs. For example, given $N = \{i_1, i_2, i_3\}$,

$$G_3 = \{(i_1, i_1), (i_1, i_2), (i_1, i_3), (i_3, i_1)\}$$

is a homogeneous directed network. Figure 3 depicts homogeneous directed network G_3 .

Here, the arc $(i_1, i_3) \in G_3$ denotes that nodes i_1 and i_3 are connected by an arc running from node i_1 to node i_3 . Note that because $(i_1, i_3) \in G_3$ there is also an arc running in the opposite direction from i_3 to i_1 . Also, note that in a homogeneous directed network, while connections have

direction, all connections are of the same type – that is, connections are homogeneous. Finally, note that in a directed network loops are allowed. For example, $(i_3, i_1) \in G_3$ and therefore in network G_3 there is an arc running from node i_1 to node i_1 .


The following definition, from (Page et al. 2005), allows for heterogeneous, multiple arcs by labeling arcs using the set A of arc types.

Definition 4 (Heterogeneous Directed Networks) A heterogeneous directed network, G , is a subset of $A \times (N \times N)$. Given any $G \subseteq A \times (N \times N)$, each ordered pair $(a, (i, i')) \in G$ consisting of an arc type and an arc is called a labeled arc in G . The collection of all labeled directed networks is denoted by $P(A \times (N \times N))$. Thus, $A \times (N \times N)$ is the set of all possible labeled arcs and a heterogeneous directed network G is simply a subset of all possible labeled arcs. For example, given $N = \{i_1, i_2, i_3\}$ and $A = \{a_1, a_2, a_3\}$ the subset


$$G_4 = \{(a_2, (i_1, i_1)), (a_1, (i_1, i_2)), (a_1, (i_1, i_3)), (a_1, (i_3, i_1)), (a_3, (i_1, i_3))\}$$

of $A \times (N \times N)$ is a heterogeneous directed network. Figure 4 depicts heterogeneous directed network G_4 .

Here, the labeled arc $(a_1, \{i_1, i_3\}) \in G_4$ denotes that nodes i_1 and i_3 are connected by an arc of type a_1 running from node i_1 to node i_3 . Note that nodes



Networks and Stability, Fig. 3 Homogeneous directed network G_3



Networks and Stability, Fig. 4 Heterogeneous directed network G_4

i_1 and i_3 are also connected by an arc of type a_3 running from node i_1 to node i_3 . Thus, in addition to having direction, connections are heterogeneous. Also, note that arc type a_1 is used three times in network G_4 : Once in describing the connection running from i_1 to i_2 , once in describing the connection running from i_1 to i_3 , and once in describing the connection running from i_3 to i_1 . Finally, note that loops are allowed. For example, $(a_2, \{i_1, i_1\}) \in G_4$ and therefore in network G_4 there is an a_2 arc running from node i_1 to node i_1 .

Remarks

1. In the terminology of graph theory (e. g., see Bollobas (1998)), a homogeneous linking network is called a graph, while a homogeneous directed network is called a directed graph.
2. The following notation is useful in describing heterogeneous directed networks. Given heterogeneous directed network $G \subseteq A \times (N \times N)$, let

$$\begin{aligned} G(a) &:= \{(i, i') \in N \times N : (a, (i, i')) \in G\}, \\ G(i, i') &:= \{a \in A : (a, (i, i')) \in G\}, \\ G^+(i) &:= \{a \in A : (a, (i, i')) \in G \text{ for some } i' \in N\}, \\ G^-(i) &:= \{a \in A : (a, (i, i')) \in G \text{ for some } i' \in N\}, \\ G^+(a, i) &:= \{i' \in N : (a, (i, i')) \in G\}, \\ G^-(a, i) &:= \{i' \in N : (a, (i, i')) \in G\}. \end{aligned}$$

For example, referring to heterogeneous directed network G_4 above (see Fig. 4),

$$\begin{aligned} G_4(a_1) &:= \{(i_1, i_2), (i_1, i_3), (i_3, i_1)\}, \\ G_4(i_1, i_3) &:= \{a_1, a_3\}, \\ G_4^+(i_1) &:= \{a_1, a_2, a_3\}, \\ G_4^-(i_1) &:= \{a_1, a_2\}, \\ G_4^+(a_2, i_1) &:= \{i_1\}, \\ G_4^-(a_2, i_1) &:= \{i_1\}. \end{aligned}$$

3. The number $|G^+(i)|$ is the number arc types leaving node i in network G , while the number $|G^-(i)|$ is the number of arc types entering node i in network G , while the number is the

indegree of node i for arc type a in network G . For directed network G_4 , we have for example,

$$\begin{aligned} |G_4^+(i_2)| &= 0 & \text{and} & & |G_4^-(i_2)| &= 1, \\ |G_4^+(a_1, i_1)| &= 2 & \text{and} & & |G_4^-(a_1, i_1)| &= 1. \end{aligned}$$

4. Rockafellar (1984), essentially defines a network G to be a nonempty subset of $A \times (N \times N)$ such that for all $a \in A$, $G(a) \subseteq \{(i, i') \in N \times N : i \neq i', \text{ and } |G(a)| \leq 1\}$.

The Feasible Set

In the abstract games of network formation we develop here, we will assume that the game is played over some feasible set of networks G . Some examples of feasible sets are: The set of all homogeneous linking networks $P(P_2(N))$ as in Jackson-Wolinsky (1996) and Jackson-van den Nouweland (2005); The set of all homogeneous directed networks $P(N \times N)$ as in Bala and Goyal (2000) and an arbitrary subset of the set of heterogeneous directed networks $P(A \times (N \times N))$ as in (2005, 2008). The following example is taken from Page and Wooders (2007) where the feasible set is taken to be the set of all club networks – a particular class of heterogeneous directed networks.

Example 1 (Club Networks) Let D be a finite set of players with typical element d and let C be a finite set of clubs or club locations with typical element c . As before, let A be a finite set of arc types. Finally, let $N = D \cup C$ be the set of nodes. We consider an abstract game of network formation played over the feasible set $G \subset P(A \times (N \times N))$ of club networks where $G \in \mathcal{G}$ if and only if G is a nonempty subset of $A \times (D \times C)$ such that (i) for all players $d \in D$, the set

$$G(d) := \{(a, c) \in A \times C : (a, (d, c)) \in G\}$$

is nonempty and (ii) for all $(a, (d, c)) \in G$, $a \in A(d, c)$. Here, $A(d, c)$ is the set of actions (represented by arc types) available to player d in club c . Given club network $G \in \mathcal{G}$, $(a, (d, c)) \in G$ means that in club network G player d is a member of club c and takes action $a \in A(d, c)$ – or in the terminology of directed networks, that in

club network G , there is an arc of type a running from node (player) d to node (club) c . Thus, in this example the feasible set is a set of bipartite directed networks.

We remark that the basic model of club formation underlying this example has a long history in the literature, going back to economies with essentially homogeneous agents modeled as games in characteristic function form (Shubik (bridge game) (Shubik 1971)) and serves as an example of several models in more recent literature on coalitional games (cf., Banerjee et al. 2001), Bogomolnaia and Jackson (2002), Diamantoudi and Xue (2003), and in economies with clubs (cf., Arnold and Wooders 2006 and Allouch and Wooders 2007). As in Konishi et al. (1998) and Demange (1994), for example, we allow “free entry” into clubs.

Paths and Circuits

A sequence of links $\{\{i, i'\}_k\}_k$ in $G \in \mathcal{G} \subseteq P(P_2(N))$ constitutes a path if each link $\{i, i'\}_k$, has one node in common with the preceding link $\{i, i'\}_{k-1}$ and the other node in common with the succeeding link $\{i, i'\}_{k+1}$. A circuit is a finite path $\{\{i, i'\}_k\}_{k=1}^h$ in G which begins at a node i and returns the same node. The length of a path is the number of links in the path.

A sequence of labeled links $\{(a, \{i, i'\})_k\}_k$ in $G \in \mathcal{G} \subseteq P(A \times P_2(N))$ constitutes a path if each labeled link $(a, \{i, i'\})_k$ has one node in common with the preceding labeled link $(a, \{i, i'\})_{k-1}$ and the other node in common with the succeeding link $(a, \{i, i'\})_{k+1}$. A circuit is a finite path $\{(a, \{i, i'\})_k\}_{k=1}^h$ in G which begins at a node i and ends at the same node. The length of a path is the number of labeled links in the path.

A sequence of arcs $\{(i, i')_k\}_k$ in $G \in \mathcal{G} \subseteq P(N \times N)$ constitutes a path if the beginning node i of arc $\{i, i'\}_k$ coincides with the ending node i' of preceding arc $\{i, i'\}_{k-1}$. A circuit is a finite path $\{(i, i')_k\}_{k=1}^h$ in G such that node i of arc $\{i, i'\}_1$ and node i' of arc $\{i, i'\}_h$ are the same node. The length of a path is the number of arcs in the path.

Finally, a sequence of labeled arcs $\{(a, (i, i'))_k\}_k$ in $G \in \mathcal{G} \subseteq P(N \times N)$ constitutes a path if the beginning node i of labeled arc $(a, (i, i'))_k$, coincides with the ending node i' of preceding arc $(a, (i, i'))_{k-1}$. A circuit is a finite path $\{(a, (i, i'))_k\}_{k=1}^h$ in G such that node i of labeled arc $(a, (i, i'))_1$ and node i' of labeled arc $(a, (i, i'))_h$ are the same node. The length of a path is the number of labeled arcs in the path.

In Fig. 5, $\{(a_1, (i_3, i_1))_1, (a_2, (i_1, i_1))_2, (a_1, (i_1, i_2))_3\}$ is a path in G_4 of length 3, while $\{(a_1, (i_3, i_1))_1, (a_2, (i_1, i_1))_2, (a_3, (i_1, i_3))_3\}$ circuit in G_4 of length 3.


Players' Preferences

For the remainder of this entry we will assume that the set of players is given by the set of nodes N . Thus, henceforth the nodes represent players in the game of network formation.

Let $\Gamma(N)$ denote the collection of all coalitions of players (i.e., nonempty subsets of N) with typical element denoted by S .

For each player $i \in N$ let $\succ_i$ be an irreflexive binary relation on $\mathcal{G}$ ($= P(P_2(N))$ or $P(N \times N)$) and write $G' \succ_i G$ if player $i \in N$ prefers network $G' \in \mathcal{G}$ to network $G \in \mathcal{G}$. Because $\succ_i$ is irreflexive, $G \not\succ_i G$ for all networks $G \in \mathcal{G}$. Coalition $S' \in \Gamma(N)$ prefers network G' to network G , written $G' \succ_{S'} G$, if $G' \not\succ_i G$ for all players $i \in S'$.

Networks and Stability,
Fig. 5 Path and circuit in
heterogeneous directed
network G_4



Note that because players' preferences $\{\succ_i\}_{i \in N}$ are irreflexive, coalitional preferences, $\{\succ_S\}_{S \in \Gamma(N)}$, are also irreflexive.

A Remark on Weak Preferences

Players are said to have weak preferences on $G (= P(P_2(N)) \text{ or } P(N \times N))$ denoted by $\succeq_i$ if $G' \succeq_i G$ means that player i either strongly prefers G' to G (denoted $G' \succ_i G$) or is indifferent between G' and G (denoted $G' \sim_i G$). If coalitional preferences are based on weak preference, then we say that coalition $S' \in P(N)$ *weakly prefers* network G' to network G , written $G' \succ_{wS'} G$, if for all players $i \in S'$, $G' \succeq_i G$ and if for at least one player $i' \in S'$, $G' \succ_{i'} G$. Note that if preferences are weak and $G' \succ_{wS'} G$ where S' consists of a single player i , so that $S' = \{i\}$, then $G' \succ_i G$. Finally, note that weak coalitional preferences $\{\succ_{wS}\}_{S \in \Gamma(N)}$ are irreflexive (i.e., $G \not\succ_{wS} G$ for all $G \in G$ and $S \in \Gamma(N)$).

Network Payoff Functions

In many applications, players' preferences are specified via real-valued network payoff functions, $\{v_i(\cdot)\}_{i \in N}$. If this is the case, then for each player $i \in N$ and each network $G \in G$, $v_i(G)$ is the payoff to player i in network G . Note that the payoff $v_i(G)$ to player i in network G depends on the entire network. Thus, the player may be affected by connections between other players even when he himself has no direct or indirect connection with those players. Intuitively, 'widespread' network externalities are allowed.

Given payoff functions $\{v_i(\cdot)\}_{i \in N}$, player i prefers network G' to network G if $v_i(G') > v_i(G)$. Coalitional preferences can then be specified by stating that coalition $S' \in \Gamma(N)$ prefers network G' to network G if $v_i(G') > v_i(G)$ for all $i \in S'$.

Preference Supernetworks

By viewing each network G in feasible set G as a node in a larger network, we can represent coalitional preferences as a heterogeneous directed network. To begin, let

$$\mathcal{P} := \{p_S : S \in \Gamma(N)\}$$

denote the set of arc labels for preference arcs (or p -arcs for short).

Definition 5 (Coalitional Preference Supernetworks, Page et al. (2005)) Given feasible set $G (= P(P_2(N)) \text{ or } P(N \times N))$, a coalitional preference supernetwork $\mathbf{P}$ is a subset of $\mathcal{P} \times (G \times G)$ such that $(p_S, (G, G'))$ is contained in $\mathbf{P}$ if and only if $G' \succ_S G$.

The Rules of Network Formation

The rules of network formation are specified via a collection of coalitional effectiveness relations $\{\rightarrow_S\}_{S \in \Gamma(N)}$ defined on the feasible set of networks $G (= P(P_2(N)) \text{ or } P(N \times N))$. Each effectiveness relation $\rightarrow_S$ represents what a coalition S can do. Thus, if $G \rightarrow_S G'$ this means that under the rules of network formation coalition $S \in \Gamma(N)$ can change network $G \in G$ to network $G' \in G$ by adding, subtracting, or replacing connections in G (where, depending on the feasible set, a *connection* is a link or an arc).

Examples of Network Formation Rules

Jackson-Wolinsky Rules (1996) (Bilateral-Unilateral Rules) Assume that the feasible set of networks G is equal to the set of homogeneous linking networks $P(P_2(N))$. Under the Jackson-Wolinsky rules of network formation (see Jackson and Wolinsky 1996),


1. adding a link from player i to player i' requires that both players i and i' agree to add the link (i.e., link addition is bilateral);
2. subtracting a link from player i to player i' requires that player i or player i' or both agree to subtract the link (i.e., link subtraction can be unilateral);
3. link addition or link subtraction takes place one link at a time.

Thus, for any pair of networks G and G' in G , if $G \rightarrow_S G'$ and $G \neq G'$, then


$$\begin{aligned} \text{either } G' &= G \cup \{i, i'\} \text{ for some } \{i, i'\} \in P_2(N) \\ &\text{and } S = \{i, i'\} \\ \text{or } G' &= G \setminus \{i, i'\} \text{ for some } \{i, i'\} \in G \\ &\text{and } S = \{i\} \text{ or } S = \{i'\} \text{ or } S = \{i, i'\}. \end{aligned}$$

To illustrate, consider Fig. 6 depicting two homogeneous linking networks G and G' .

Networks and Stability,
Fig. 6 (a) Network G . (b)
 Network G'



Networks and Stability,
Fig. 7 (a) Network G . (b)
 Network G''



Observe that

$$G' = G \cup \{i_2, i_3\} \text{ and } G = G' \setminus \{i_2, i_3\}.$$

Under the effectiveness relations implied by the Jackson-Wolinsky rules, for networks G and G' we have

$$G \xrightarrow{\{i_2, i_3\}} G', \quad G' \xrightarrow{\{i_2, i_3\}} G, \quad G' \xrightarrow{\{i_2\}} G, \quad G' \xrightarrow{\{i_3\}} G.$$

Jackson-van den Nouweland Rules (Jackson and van den Nouweland 2005) (Bilateral-Unilateral Rules) Again assume that the feasible set of networks G is equal to the set of homogeneous linking networks $P(P_2(N))$. Under the Jackson-van den Nouweland rules of network formation,

1. adding a link from player i to player i' requires that both players i and i' agree to add the link (i.e., link addition is bilateral);
2. subtracting a link from player i to player i' requires that player i or player i' or both agree to subtract the link (i.e., link subtraction can be unilateral).

Thus, the Jackson-van den Nouweland rules are the Jackson-Wolinsky rules without the one-link-at-a-time restriction. Note that if link

addition is bilateral and link subtraction is unilateral (i.e., if rules (i) and (ii) hold), then $G \rightarrow_S G''$ and $G \neq G''$ implies that

- (i) if $\{i, i'\} \in G''$ and $\{i, i'\} \notin G$, then $\{i, i'\} \subseteq S$;

and

- (ii) if $\{i, i'\} \notin G''$ and $\{i, i'\} \in G$, then $\{i, i'\} \cap S \neq \emptyset$.

To illustrate, consider Fig. 7 depicting two homogeneous linking networks G and G'' .

Observe that

$$G'' = (G \setminus \{i_1, i_3\}) \cup \{i_2, i_3\}$$

and

$$G = (G'' \setminus \{i_2, i_3\}) \cup \{i_1, i_3\}.$$

Under the effectiveness relations implied by the Jackson-van den Nouweland rules, for networks G and G'' we have

$$G \xrightarrow{\{i_1, i_2, i_3\}} G'', \quad G \xrightarrow{\{i_2, i_3\}} G'', \quad G'' \xrightarrow{\{i_1, i_3\}} G, \\ G'' \xrightarrow{\{i_1, i_2, i_3\}} G.$$

Note that under the one-link-at-a-time restriction, it is not possible under the Jackson-Wolinsky

rules to move directly from network G to network G'' or directly from network G'' to network G (i.e., G and G'' are *not* related under the effectiveness relations $\{\rightarrow_S\}_{S \in \Gamma(N)}$). Instead, under the Jackson-Wolinsky rules, the change from G to G'' or from G'' to G requires two moves. For example,

first $G \xrightarrow[\{i_2, i_3\}]{} G'$, and then, $G' \xrightarrow[\{i_3\}]{} G''$ or $G' \xrightarrow[\{i_1\}]{} G''$;

or

first $G'' \xrightarrow[\{i_1, i_3\}]{} G'$, and then, $G' \xrightarrow[\{i_3\}]{} G$ or $G' \xrightarrow[\{i_2\}]{} G$.

Bala-Goyal Rules (Bala and Goyal 2000)

(Noncooperative Rules – Unilateral-Unilateral Rules)

Now assume that the feasible set of networks is equal to the set of homogeneous directed networks $P(N \times N)$. Translating Bala and Goyal rules into our notation and terminology,

1. adding an arc from player i to player i' requires only that player i agree to add the arc (i.e., arc addition is unilateral and can be carried out only by the initiator, player i);
2. subtracting an arc from player i to player i' requires only that player i agree to subtract the arc (i.e., arc subtraction is unilateral and can be carried out only by the initiator, player i);
3. $G \rightarrow_S G'$ implies that $|S| = 1$ (i.e., only network changes brought about by individual players are allowed).

We shall also refer to rules (i)-(iii) as noncooperative. Note that a player i can add or subtract an

arc to player i' without regard to the preferences of player i' and can add and/or subtract arcs to several players simultaneously and can do so without regard to those players' preferences. Thus in general under noncooperative rules, effectiveness relations display a type of symmetry, and in particular, if $G \rightarrow_{\{i\}} G'$ then $G' \rightarrow_{\{i\}} G$.

To illustrate, consider Fig. 8 depicting three homogeneous directed networks G , G' , and G'' .

Under the effectiveness relations implied by noncooperative rules for networks G and G' in Fig. 8 we have

$$G \xrightarrow[\{i_1\}]{} G', \quad G' \xrightarrow[\{i_1\}]{} G.$$

Note that under noncooperative rules, networks G and G'' in Fig. 8 are *not* related under the effectiveness relations $\{\rightarrow_{\{i\}}\}_{i \in N}$. However, under the noncooperative rules we have, for example, the following effectiveness relations

$$G \xrightarrow[\{i_1\}]{} G', \quad G' \xrightarrow[\{i_2\}]{} G''$$

and


$$G'' \xrightarrow[\{i_2\}]{} G', \quad G' \xrightarrow[\{i_1\}]{} G.$$

Rules Supernetworks

Again by viewing each network G in feasible set $\mathcal{G}$ as a node in a larger network, we can represent the rules of network formation as a heterogeneous directed network. To begin, let

$$\mathcal{M} := \{m_S : S \in \Gamma(N)\}$$

denote the set of arc labels for move arcs (or m -arcs for short).



Networks and Stability, Fig. 8 (a) Network G . (b) Network G' . (c) Network G''

Definition 6 (Rules Supernetworks, Page et al. (2005)) Given feasible set $G (= P(P_2(N))$ or $P(N \times N)$), a rules supernetwork $\mathbf{R}_\rho$ is a subset of $\mathcal{M} \times (G \times G)$ such that $(m_{S'}, (G, G'))$ is contained in $\mathbf{R}_\rho$ if and only if $G \rightarrow_{S'} G'$, where ρ denotes the name of the network formation rules in force. We shall adopt the convention that $\rho = jw$ if the rules are Jackson-Wolinsky, $\rho = jn$ if the rules are Jackson-van den Nouweland, and $\rho = bg$ if the rules are Bala-Goyal or noncooperative.

Supernetworks

Given feasible set $G (= P(P_2(N))$ or $P(N \times N)$), coalitional preferences $\{\succ_S\}_{S \in \Gamma(N)}$ coalitional effectiveness relations $\{\rightarrow_S\}_{S \in \Gamma(N)}$ can be represented by a heterogeneous directed network called a supernetwork (see Page and Kamat 2005; Page et al. 2005). In particular, given preference supernetwork $\mathbf{P}$ and rules supernetwork $\mathbf{R}_\rho$ the corresponding supernetwork is given by

$$\mathbf{G}_\rho := \mathbf{P} \cup \mathbf{R}_\rho$$

Letting $\mathcal{A} := \mathcal{P} \cup \mathcal{M}$ (i.e., the union of all preference arcs and move arcs), then

$$\mathbf{G}_\rho \subset \mathcal{A} \times (G \times G).$$

Dominance Relations

Direct Dominance

Given feasible set $G (= P(P_2(N))$ or $P(N \times N)$), coalitional preferences $\{\succ_S\}_{S \in \Gamma(N)}$ and coalitional effectiveness relations $\{\rightarrow_S\}_{S \in \Gamma(N)}$, network $G' \in G$ *directly dominates* network $G \in G$, written $G' \triangleright G$, if for some coalition $S' \in \Gamma(N)$,

$$G \prec_{S'} G' \quad \text{and} \quad G \rightarrow_{S'} G'$$

Thus, network G' directly dominates network G if some coalition S' prefers G' to G and if under the rules of network formation coalition S' has the power to change G to G' .

Note that direct dominance is irreflexive but not in general transitive. Also note that if $\mathbf{G}_\rho$ is the supernetwork, then $G' \triangleright G$ if and only if

$(p_{S'}, (G, G')) \in \mathbf{G}_\rho$ and $(m_{S'}, (G, G')) \in \mathbf{G}_\rho$ for some coalition S' .

Indirect Dominance

Given feasible set $G (= P(P_2(N))$ or $P(N \times N)$), coalitional preferences $\{\succ_S\}_{S \in \Gamma(N)}$ and coalitional effectiveness relations $\{\rightarrow_S\}_{S \in \Gamma(N)}$ network $G' \in G$ *indirectly dominates* network $G \in G$, written $G' \triangleright \triangleright G$, if there is a *finite* sequence of networks,

$$G_0, G_1, \dots, G_h,$$

with $G = G_0$, $G' = G_h$, and $G_k \in G$ for $k = 0, 1, \dots, h$, and a corresponding sequence of coalitions,

$$S_1, S_2, \dots, S_h,$$

such that for $k = 1, 2, \dots, h$

$$G_{k-1} \xrightarrow{S_k} G_k, \quad \text{and} \quad G_{k-1} \prec_{S_k} G_k.$$

Note that if network G' indirectly dominates network G (i.e., if $G' \triangleright \triangleright G$), then what matters to the initially deviating coalition S_1 , as well as all the coalitions along the way, is that the ultimate network outcome $G' = G_h$ be preferred. Thus, for example, the initially deviating coalition S_1 will not be deterred from changing network G_0 to network G_1 even if network G_1 is not preferred to network $G = G_0$, as long as the ultimate network outcome $G' = G_h$ is preferred to G_0 that is, as long as $G_0 \prec_{S_1} G_h$. Finally, note that indirect dominance is irreflexive but not in general transitive.

In order to capture the idea of farsightedness in strategic behavior, Chwe (1994) analyzed abstract games equipped with indirect dominance relations in great detail, introducing the equilibrium notions of consistency and largest consistent set. The basic idea of indirect dominance goes back to the work of Guilbaud (1949) and Harsanyi (1974).

Given the supernetwork representation of preferences and rules, $\mathbf{G}_\rho$, we can write, $G' \triangleright \triangleright G$ if there is a *finite* sequence of networks,

$$G_0, G_1, \dots, G_h,$$

with $G = G_0$, $G' = G_h$, and $G_k \in G$ for $k = 0, 1, \dots, h$, and a corresponding sequence of coalitions,

$$S_1, S_2, \dots, S_h,$$

such that for $k = 1, 2, \dots, h$

$$(m_{S_k}, (G_{k-1}, G_k)) \in \mathbf{G}_\rho,$$

and

$$(p_{S_k}, (G_{k-1}, G_k)) \in \mathbf{G}_\rho.$$

Path Dominance

Any irreflexive dominance relation $>$ on $G (= P(P_2(N)) \text{ or } P(N \times N))$ – for example, direct or indirect – induces a path dominance relation on the set of networks (sometimes referred to as the transitive closure of $>$). In particular, corresponding to dominance relation $>$ on networks G there is a corresponding path dominance relation $\geq_p$ on G specified as follows: Network $G' \in G$ path dominates network $G \in G$ with respect to $>$ (i.e., with respect to the underlying dominance relation $>$), written $G' \geq_p G$, if $G' = G$ or if there exists a *finite* sequence of networks $\{G_k\}_{k=0}^h$ in G with $G_h = G'$ and $G_0 = G$ such that for $k = 1, 2, \dots, h$

$$G_k > G_{k-1}.$$


We refer to such a finite sequence of networks as a *finite domination path* and we say network G' is $>$ – *reachable* from network G if there exists a finite domination path from G to G' . Thus,

$$G' \geq_p G \text{ if and only if } \begin{cases} G' \text{ is } > \text{ – reachable from } G, \text{ or} \\ G' = G \end{cases}$$

Note that, even though the underlying dominance relation $>$ is irreflexive and intransitive or transitive, the induced path dominance relation $\geq_p$ on G is both reflexive ($G \geq_p G$) and transitive ($G' \geq_p G$ and $G'' \geq_p G'$ implies that $G'' \geq_p G$).

$>$ –Supernetworks

Let $>$ denote the irreflexive dominance relation on G . It is often useful to represent $>$ as a *homogeneous* directed network, $\mathbf{D}_>$, where $\mathbf{D}_>$ is a



Networks and Stability, Fig. 9 $>$ –Supernetwork $\mathbf{D}_>$

subset of $G \times G$ and where $>$ –arc $(G, G') \in \mathbf{D}_>$ if and only if $G' > G$ (i.e., if and only if $G' >$ –dominates G). We call such a homogeneous directed network (or directed graph) a $>$ –supernetwork. For example, suppose $G = \{G_1, G_2, G_3, \dots, G_7\}$ and suppose the dominance relation $>$ on G is a direct dominance relation and has the supernetwork representation given in Fig. 9.

Note that network G_5 is $>$ –reachable through $\mathbf{D}_>$ from network G_1 by the domination path given by the $>$ –arc sequence

$$\{(G_1, G_6)_1, (G_6, G_2)_2, (G_2, G_3)_3, (G_3, G_4)_4, (G_4, G_5)_5\}.$$

Thus, G_5 path dominates G_1 . Note that network G_2 is $>$ –reachable from network G_2 by the domination circuit given by the $>$ –arc sequence

$$\{(G_2, G_3)_1, (G_3, G_4)_2, (G_4, G_5)_3, (G_5, G_2)_4\}.$$

and that network G_3 is $>$ –reachable from network G_3 by two domination circuits given by the $>$ –arc sequences

$$\{(G_3, G_4)_1, (G_4, G_5)_2, (G_5, G_2)_3, (G_2, G_3)_4\}$$

and

$$\{(G_3, G_4)_1, (G_4, G_3)_2\}.$$

Because networks G_2 and G_5 are on the same circuit, G_5 is $>$ -reachable from G_2 and G_2 is $>$ -reachable from G_5 . Thus, G_5 path dominates G_2 (i.e., $G_5 \geq_p G_2$) and G_2 path dominates G_5 (i.e., $G_2 \geq_p G_5$). The same cannot be said of networks G_1 and G_5 . In particular, while $G_5 \geq_p G_1$ it is not true that $G_1 \geq_p G_5$ because G_1 is not $>$ -reachable from G_5 . Finally, note that network G_0 is isolated in $\mathbf{D}_>$. In particular, G_0 is *not* reachable through $\mathbf{D}_>$ from any network in G and no network in G is reachable through $\mathbf{D}_>$ from G . In general, a network $G \in G$ is isolated if there does not exist a network $G' \in G$ with $G' \geq_p G$ or $G \geq_p G'$.

Note that if the direct dominance relation with $>$ -supernetwork depicted in Fig. 9 has underlying coalitional preferences $\{\succ_S\}_{S \in \Gamma(N)}$ and coalitional effectiveness relations $\{\rightarrow_S\}_{S \in \Gamma(N)}$ then the $>$ -arc from network G_3 to network G_4 in Fig. 9 means that for some coalition S , G_4 is preferred to G_3 and more importantly, that coalition S has the power to change network G_3 to network G_4 . Thus, $G_3 \prec_{sS} G_4$ and $G_3 \rightarrow_S G_4$. But because there is a $>$ -arc in the opposite direction, from network G_4 to network G_3 , G_3 also directly dominates G_4 . Thus for some coalition S' disjoint from coalition S ($S' \cap S = \emptyset$), $G_4 \prec_{sS'} G_3$ and $G_4 \rightarrow_{S'} G_3$ – Finally, note that if coalitional preferences over networks are weak (i.e., are based on weak preferences), then the statement, ‘for some coalition S' disjoint from coalition S can be weakened to for some coalition S' not equal to coalition S . With this weakening, the requirement that the intersection of S and S' be empty is no longer needed.

Abstract Games of Network Formation and Stability

An abstract game of network formation consists of a feasible set of networks equipped with a dominance relation. We shall consider two classes of games: (i) games where the feasible set of networks $G (= P(P_2(N)) \text{ or } P(N \times N))$ is equipped with an irreflexive dominance relation $>$, either

direct or indirect (i.e., $>$ is equal to $\triangleright$ or $\triangleright\triangleright$), induced by coalitional preferences and network formation rules; and (ii) games where the feasible set of networks is equipped with a path dominance relation $\geq_p$ induced by such an irreflexive dominance relation

Network Formation Games with Respect to Irreflexive Dominance

In this section we consider the abstract game with respect to irreflexive dominance given by the pair

$$(G, >).$$

Throughout this section we will assume that primitives are represented by supernetwork $\mathbf{G}_\rho := \mathbf{P} \cup \mathbf{R}_\rho$ (where ρ is equal to jw , jn , or bg) and $>$ -supernetwork $\mathbf{D}_>$ (where $>$ is equal to $\triangleright$ or $\triangleright\triangleright$), and that the feasible set of networks G is equal to the set of homogeneous linking networks $P(P_2(N))$ or the set of homogeneous directed networks $P(N \times N)$.

Quasi-Stability and Stability

We define the $>$ -distance from G_0 to G_1 in $\mathbf{D}_>$ to be the length of the *shortest* $>$ -path from G_0 to G_1 if G_1 is $>$ -reachable from G_0 in $\mathbf{D}_>$, and $+\infty$ if G_1 is not reachable from G_0 in $\mathbf{D}_>$. We denote the distance from G_0 to G_1 in $\mathbf{D}_>$ by $d_{\mathbf{D}_>}(G_0, G_1)$. Thus,

$$d_{\mathbf{D}_>}(G_0, G_1) = \begin{cases} \text{length of shortest } > \text{-path} \\ \text{from } G_0 \text{ to } G_1 \text{ in } \mathbf{D}_>, \\ \text{if } G_1 \text{ is reachable from } G_0, \\ +\infty, G_1 \text{ not reachable from} \\ G_0 \text{ in } \mathbf{D}_>. \end{cases}$$

The following are network renditions of quasi-stable and stable sets.

Definition 7 (Quasi-Stable Sets and Stable Sets, (Berge 2001; Chvatal and Lovasz 1972))

1. A subset Q of networks in G is said to be quasi-stable for network formation game $(G, >)$ if,
 1. Q is internally stable, that is, $d_{\mathbf{D}_>}(G_0, G_1) \geq 2$, whenever G_0 and G_1 are in Q , with $G_0 \neq G_1$, and

2. Q is externally quasi-stable, that is, given any $G_0 \notin Q$, there exists $G_1 \in Q$ with $d_{\mathbf{D}_>}(G_0, G_1) \leq 2$.
2. A subset S of networks in G is said to be stable for network formation game $(G, >)$ if,
 1. S is internally stable and
 2. S is externally stable, that is, given any $G_0 \notin S$, there exists $G_1 \in S$ with $d_{\mathbf{D}_>}(G_0, G_1) \leq 1$.

Thus, if Q is externally quasi-stable, a path of length at most 2 is required to get from any network outside of Q to a network in Q , whereas, if Q is externally stable a path of length at most 1 is required.

Letting

$$P_{>}(G_0) := \{G \in G : G > G_0\},$$

an alternative way to write part (2) of the definition above is as follows:

(2)' A subset S of networks in G is said to be stable if

1. (internal stability) $G \in S$ implies that $P_{>}(G) \cap S = \emptyset$ and
2. (external stability) $G \notin S$ implies that $P_{>}(G) \cap S \neq \emptyset$.

If S is stable (or quasi-stable), then it is automatically nonempty. Note that a stable set S is simply a von Neuman-Morgenstern stable set with respect to the dominance relation $>$ defined on G . Also, note that if S is stable then it is automatically quasi-stable.

We now state a remarkably simple result on the existence of quasi-stable sets. This result is a network rendition of a general result due to Chvatal and Lovasz (1972) on the existence of quasi-stable sets in directed graphs (here the directed graph is the $>$ -supernetwork $\mathbf{D}_>$; also see, Galeana-Sanchez and Xueliang (1998)).

Theorem 1 (Existence of Quasi-Stable Sets for Network Formation Games, Page and Kamat (2005)) There exists a quasi-stable set Q for network formation game $(G, >)$.

In fact, it follows from the Theorem due to Chvatal and Lovasz (1972) that any finite set Z equipped with an irreflexive binary relation $<$ has a $<$ -quasi-stable set. Moreover, if the relation $<$ is transitive, then any $<$ -quasi-stable set is $<$ -stable.

Next we state two results on the existence of stable sets.

Theorem 2 (Existence of Stable Sets for Network Formation Games, Page and Kamat (2005))

1. If $\mathbf{D}_>$ contains no $>$ -circuits, then there exists a unique stable S for $(G, >)$.
2. If $\mathbf{D}_>$ contains no $>$ -circuits of odd length, then there exists a stable S for $(G, >)$.

Part (1) of Theorem 2 is an immediate consequence of a 1958 result due to Berge (see Theorem 4, p. 48 in Berge (2001)). Part (2) of Theorem 2 is a supernetwork version of the classical result due to Richardson (1953). A $>$ -circuit in supernetwork $\mathbf{D}_>$ is said to be of odd length if there is an odd number of connections in the circuit.

Farsighted Consistency

Chwe (1994) in an influential paper introduced the notion of farsighted consistency in an abstract game. The following is a definition of farsighted consistency for abstract games of network formation.

Definition 8 (Farsighted Consistency, Page and Kamat (2005)) A subset F of networks in G is said to be farsightedly consistent for network formation game $(G, \triangleright)$ if

for all $G_0 \in F$,

$$(m_{S_1}, (G_0, G_1)) \in \mathbf{G}_\rho \text{ for some } G_1 \in G$$

and some coalition S_1 , implies that

$$\text{there exists } G_2 \in F$$

with $G_2 = G_1$ or $G_2 \triangleright G_1$ such that,

$$(p_{S_1}, (G_0, G_1)) \notin \mathbf{G}_\rho.$$

In words, a subset of directed networks F is said to be farsightedly consistent if given any

network $G_0 \in F$ and any m_{S_1} -deviation to network $G_1 \in G$ by coalition S_1 (via adding, subtracting, or replacing arcs in accordance with $\mathbf{R}_\rho$ there exists further deviations leading to some network $G_2 \in F$ where the initially deviating coalition S_1 is not better off – and possibly worse off. A network $G \in G$ is said to be farsightedly consistent if $G \in F$ where F is a farsightedly consistent set.

If (i) $G = P(P_2(N))$, (ii) coalitional preferences are weak, denoted by $\{\succ_{wS}\}_{S \in \Gamma(N)}$, so that indirect dominance is weak (denoted by $\triangleright_w$ and defined in the obvious way), and (iii) coalitional effectiveness relations are determined by Jackson-Wolinsky rules, then the notion of farsighted consistency above (essentially due to Chwe (1994)) is closely related to the notion of pairwise farsighted stability introduced in Herrings et al. (2006).

For any game $(G, \triangleright)$, there can be many farsightedly consistent sets. We shall denote by F^* the *largest farsightedly consistent set*. Thus, if F is farsightedly consistent, then $F \subseteq F^*$. Unlike quasi-stable sets and stable sets where existence implies nonemptiness, in considering farsightedly consistent sets two critical questions arise: (i) does there exist a largest farsightedly consistent set of networks for $(G, \triangleright)$, and (ii) is it nonempty? Our next result provides a positive answer to both questions.

Theorem 3 (Existence, Uniqueness, and Nonemptiness of F^* , Page et al. (2005)) There exists a unique, nonempty largest farsightedly consistent set F^* for network formation game $(G, \triangleright)$. Moreover, F^* is externally stable; that is, if network G is not contained in F^* , then there exists a network G' contained in F^* that indirectly dominates G (i.e., $G' \triangleright \triangleright G$).

The method of proving existence and uniqueness is a straightforward, supernetwork rendition of Chwe's (1994) method and is similar to the method introduced by Roth (1975, 1977), Page and Kamat (2005) provide an alternative proof (to that of Chwe and of (Page et al. 2005)) of the nonemptiness and external stability of the largest consistent set (with respect to indirect dominance). In particular, Page and Kamat modify the indirect dominance relation so as to make it transitive as well as irreflexive. They then show that

the unique stable set with respect to path dominance induced by this new transitive indirect dominance relation is contained in the largest farsightedly consistent set – and in this way show that the largest farsightedly consistent set is nonempty and externally stable.

Network Formation Games with Respect to Path Dominance

In this section we consider the network formation game with respect to path dominance given by the pair

$$(G, \geq_p).$$

Throughout this section we will assume that the underlying primitives,

$$(G, \{\succ_S\}, \{\rightarrow_S\}, \triangleright)_{S \in \Gamma(N)},$$

are such that G is equal to the set of homogeneous linking networks $P(P_2(N))$ or the set of homogeneous directed networks $P(N \times N)$, that the dominance relation $\triangleright$ on G is given by either direct dominance $\triangleright$ or indirect dominance $\triangleright \triangleright$, and that primitives are represented by supernetwork $\mathbf{G}_\rho := \mathbf{P} \cup \mathbf{R}_\rho$ (with ρ equal to jw , jn , or bg) and $\triangleright$ -supernetwork $\mathbf{D}_\triangleright$.

We will present three notions of stability introduced in (Page and Wooders 2005, 2008) for abstract games of network formation with respect to path dominance over heterogeneous directed networks: (i) strategic basins of attraction, (ii) path dominance stable sets, and (iii) the path dominance core.

Preliminaries

Networks Without Descendants If $G_1 \geq_p G_0$ and $G_0 \geq_p G_1$, networks G_1 and G_0 are *equivalent*, written $G_1 \equiv_p G_0$. If networks G_1 and G_0 are equivalent then either networks G_1 and G_0 coincide or G_1 and G_0 are on the same circuit (see Fig. 9 above for a picture of a circuit). If $G_1 \geq_p G_0$ but G_1 and G_0 are not equivalent (i.e., not $G_1 \equiv_p G_0$), then network G_1 is a *descendant* of network G_0 and we write

$$G_1 >_p G_0.$$

Referring to Fig. 9, observe that network G_5 is a descendant of network G_1 , that is, $G_5 >_p G_1$.

Network $G' \in G$ has no descendants in G if for any network $G \in G$

$$G \geq_p G' \text{ implies that } G \equiv_p G'.$$

Thus, if G' has no descendants then $G \geq_p G'$ implies that G and G' coincide or lie on the same circuit. Note that any isolated network is by definition a network without descendants (e. g., network G_0 in Fig. 9).

In attempting to identify and characterize stable homogeneous networks, networks *without descendants* are of particular interest. Here is our main result concerning networks without descendants.

Theorem 4 (All Path Dominance Network Formation Games Have Networks Without Descendants, Page and Wooders (2005, 2008)) In network formation game $(G, \geq_p)$ every network $G \in G$ is path dominated by a network $G' \in G$ without descendants (i.e., $G' \geq_p G$ and G' has no descendants).

By Theorem 4, in network formation game $(G, \geq_p)$, corresponding to any network $G \in G$ there is a network $G' \in G$ without descendants which is $>$ -reachable from G . Thus, in any network formation game the set of networks without descendants is nonempty. Referring to Fig. 9, the set of networks without descendants is given by

$$\{G_0, G_2, G_3, G_4, G_5, G_7\}.$$

We shall denote by Z the set of networks without descendants.

Basins of Attraction

Stated loosely, a basin of attraction is a set of *equivalent* networks to which the strategic network formation process represented by the game might tend and from which there is no escape. Formally, we have the following definition.

Definition 9 (Basin of Attraction, Page and Wooders (2005, 2008)) A set of networks $A \subseteq G$ is said to be a basin of attraction for $(G, \geq_p)$ if

1. the networks contained in A are equivalent (i.e., for all G' and G in A , $G' \equiv_p G$) and for no set A' having A as a strict subset is this true that all the networks in A' are equivalent, and
2. no network in A has descendants (i.e., there does not exist a network $G' \in G$ such that $G' >_p G$ for some $G \in A$).

A is a strict subset of A' if $A \subset A'$ and $A' \setminus A \neq \emptyset$.

As the following characterization result shows, there is a very close connection between networks without descendants and basins of attraction.

Theorem 5 (A Characterization of Basins of Attraction, Page and Wooders (2005, 2008))

Let A be a subset of networks in G . The following statements are equivalent:

1. A is a basin of attraction for $(G, \geq_p)$.
2. There exists a network without descendants, $G \in Z$, such that

$$A = \{G' \in Z : G' \equiv_p G\}.$$

In light of Theorem 5, we conclude that in any network formation game $(G, \geq_p)$, G contains a *unique*, finite, disjoint collection of basins of attraction, say $\{A_1, A_2, \dots, A_m\}$, where for each $k = 1, 2, \dots, m$ ($m \geq 1$)

$$A_k = A_G := \{G' \in Z : G' \equiv_p G\}$$

for some network $G \in Z$. Note that for networks G' and G in Z such that $G' \equiv_p G$, $A_{G'} = A_G$ (i.e. the basins of attraction $A_{G'}$ and A_G coincide). Also, note that if network $G \in G$ is isolated, then $G \in Z$ and

$$A_G := \{G' \in Z : G' \equiv_p G\} = \{G\}$$

is, by definition, a basin of attraction – but a very uninteresting one.

Example 2 (Basins of Attraction) In Fig. 9 above the set of networks without descendants is given by

$$Z = \{G_0, G_2, G_3, G_4, G_5, G_7\}.$$

Even though there are six networks without descendants, because networks G_2 , G_3 , G_4 , and G_5 are equivalent, there are only three basins of attraction:

$$A_1 = \{G_0\}, A_2 = \{G_2, G_3, G_4, G_5\}, \text{ and } A_3 = \{G_7\}.$$

Moreover, because G_2 , G_3 , G_4 , and G_5 are equivalent,

$$A_{G_2} = A_{G_3} = A_{G_4} = A_{G_5} = \{G_2, G_3, G_4, G_5\}.$$

Stable Sets with Respect to Path Dominance

The formal definition of a $\geq_p$ – stable set is as follows.

Definition 10 (Stable Sets with Respect to Path Dominance, Page and Wooders (2005, 2008))

A subset V of networks in G is said to be a stable set for $(G, \geq_p)$ if

1. (internal $\geq_p$ -stability) whenever G_0 and G_1 are in V , with $G_0 \neq G_1$, then neither $G_1 \geq_p G_0$ nor $G_0 \geq_p G_1$ hold, and
2. (external $\geq_p$ -stability) for any $G_0 \notin V$ there exists $G_1 \in V$ such that $G_1 \geq_p G_0$.

In other words, a nonempty subset of networks V is a stable set for $(G, \geq_p)$ if G_0 and G_1 are in V , with $G_0 \neq G_1$, then G_1 is not $>$ -reachable from G_0 , nor is $G_0 >$ -reachable from G_1 , and if $G_0 \notin V$, then there exists $G_1 \in V$ reachable from G_0 .

We now have our main results on the existence, construction, and cardinality of stable sets. These results can be viewed as variations on some classical results from graph theory applied to network formation games (e.g., see Berge 2001, Chap. 2).

Theorem 6 (Stable Sets: Existence, Construction, and Cardinality, Page and Wooders (2005, 2008)) Without loss of generality assume that $(G, \geq_p)$ has basins of attraction given by

$$\{A_1, A_2, \dots, A_m\},$$

where basin of attraction A_k contains $|A_k|$ many networks (i.e., $|A_k|$ is the cardinality of A_k). Then the following statements are true:

1. $V \subseteq G$ is a stable set for $(G, \geq_p)$ if and only if V is constructed by choosing one network from each basin of attraction, that is, if and only if V is of the form

$$V = \{G_1, G_2, \dots, G_m\},$$

where $G_k \in A_k$ for $k = 1, 2, \dots, m$.

2. $(G, \geq_p)$ possesses

$$|A_1| \cdot |A_2| \cdot \dots \cdot |A_m| := M$$

many stable sets and each stable set, V_q , $q = 1, 2, \dots, M$, has cardinality

$$|V_q| = |\{A_1, A_2, \dots, A_m\}| = m.$$

Example 3 (Basins of Attraction and Stable Sets) Referring to Fig. 9, it follows from Theorem 6 that because

$$|A_1| \cdot |A_2| \cdot |A_3| = 1 \cdot 4 \cdot 1 = 4,$$

the network formation game $(G, \geq_p)$ has 4 stable sets, each with cardinality 3. By examining Fig. 9 in light of Theorem 6, we see that the stable sets for $(G, \geq_p)$ are given by

$$\begin{aligned} V_1 &= \{G_0, G_2, G_7\}, & V_2 &= \{G_0, G_3, G_7\}, \\ V_3 &= \{G_0, G_4, G_7\}, & V_4 &= \{G_0, G_5, G_7\}. \end{aligned}$$

It should be noted that by equipping the abstract network formation game with the path dominance relation rather than the original dominance relation, we avoid the famous Lucas (1968) example of a game with no stable set.

The Path Dominance Core

Definition 11 (The Path Dominance Core, Page and Wooders (2005, 2008)) A network $G \in G$ is

contained in the path dominance core C of network formation game $(G, \geq_p)$ if and only if there does not exist a network $G' \in G$, $G' \neq G$, such that $G' \geq_p G$.

Our next results give necessary and sufficient conditions for the path dominance core of a network formation game over homogeneous networks to be nonempty, as well as a recipe for constructing the path dominance core.

Theorem 7 (Path Dominance Core: Non-emptiness and Construction, Page and Wooders (2005, 2008)) Without loss of generality assume that $(G, \geq_p)$ has basins of attraction given by

$$\{A_1, A_2, \dots, A_m\},$$

where basin of attraction A_k contains $|A_k|$ many networks. Then the following statements are true:

1. $(G, \geq_p)$ has a nonempty path dominance core if and only if there exists a basin of attraction containing a single network, that is, if and only if for some basin of attraction A_k , $|A_k| = 1$.
2. Let

$$\{A_{k_1}, A_{k_2}, \dots, A_{k_n}\} \subseteq \{A_1, A_2, \dots, A_m\},$$


be the subset of basins of attraction containing all basins having cardinality 1. Then the path dominance core C of $(G, \geq_p)$ is given by

$$C = \{G_{k_1}, G_{k_2}, \dots, G_{k_n}\},$$

where $G_{k_i} \in A_{k_i}$, for $i = 1, 2, \dots, n$.

If coalitional preferences over networks are based on weak preferences, that is, if coalitional preferences are given by $\{\succ_{wS}\}_{S \in \Gamma(N)}$, then the corresponding path dominance core – the weak path dominance core – is contained in the path dominance core based on strong preference relations.

Example 4 (Basins of Attraction and the Path Dominance Core, Page and Wooders (2005, 2008)) It follows from Theorem 7 that the path dominance core of the network formation game $(G, \geq_p)$ with feasible set $G = \{G_0, G_1, \dots, G_7\}$ and path dominance relation $\geq_p$ induced by the dominance relation depicted in Fig. 9 is



Networks and Stability, Fig. 10 A different $>$ -supernetwork $D_{>}$

$$C = \{G_0, G_7\}.$$

Suppose that the $>$ -supernetwork corresponding to the direct dominance relation $>$ on $G = \{G_0, G_1, \dots, G_7\}$ is instead depicted by Fig. 10.

Now the network formation game $(G, \geq_p)$ has 3 circuits and 1 basin of attraction, $A = \{G_2, G_3, G_4, G_5\}$. Because $|A_1| = 4$, by Theorem 7 the path dominance core of $(G, \geq_p)$ is empty. By Theorem 6, $(G, \geq_p)$ has 4 stable sets each containing 1 network. These stable sets are given by

$$V_1 = \{G_2\}, V_2 = \{G_3\}, V_3 = \{G_4\}, V_4 = \{G_5\}.$$

The Path Dominance Core and Constrained Pareto Efficiency

We say that a network $G \in G$ is constrained Pareto efficient for game $(G, \geq_p)$ if and only if there does not exist another network $G' \in G$ such that (i) some coalition S can change network G to network G' (that is, $G \rightarrow_S G'$ for some coalition $S \in \Gamma(N)$) and (ii) G' is preferred by all players (that is, $G <_i G'$ for all players $i \in N$). Letting E denote the set of all constrained Pareto efficient networks, it is easy to see that the path dominance core C of network formation game $(G, \geq_p)$ is a subset of E , that is, $C \subseteq E$.

Under the classical notion of Pareto efficiency, a network G is said to be Pareto efficient if and only if there does not exist another network G' such that $G \prec_i G'$ for all players $i \in N$, regardless of whether or not some coalition S can change network G to network G' . Letting PE denote the set of all classically Pareto efficient networks, it is easy to see that $PE \subseteq E$. Note, however, that if under the rules of network formation, any network G can be changed to any other network G' via the actions of *some* coalition S , then the notions of constrained Pareto efficiency and classical Pareto efficiency are equivalent. Thus, if the collection of coalitional effectiveness relations $\{\rightarrow_S\}_{S \in \Gamma(N)}$ on G is complete, that is, if for any pair of networks G and G' in G , $G \rightarrow_S G'$ for some coalition $S \in \Gamma(N)$, then $PE = E$, and we have $C \subseteq PE = E$.

Strong Stability, Pairwise Stability, Nash Stability, and Farsighted Consistency

In this section we continue our discussion of the network formation game with respect to path dominance given by the pair

$$(G, \geq_p).$$

Throughout this section we will continue to assume that the underlying primitives,

$$(G, \{\succ_S\}, \{\rightarrow_S\}, \triangleright)_{S \in \Gamma(N)},$$

are such that G is equal to the set of homogeneous linking networks $P(P_2(N))$ or the set of homogeneous directed networks $P(N \times N)$, that the dominance relation $\triangleright$ on G is given by either direct dominance $\triangleright$ or indirect dominance $\triangleright^*$, and that primitives are represented by supernetwork $\mathbf{G}_\rho := \mathbf{P} \cup \mathbf{R}_\rho$ (with ρ equal to jw , jn , or bg) and $\triangleright$ -supernetwork $\mathbf{D}_\triangleright$.

We will complete our unification of stability results for network formation games by concluding that, depending on how we further specialize the primitives underlying the game $(G, \geq_p)$, the path dominance core is equal to the set of pairwise

stable networks (Jackson and Wolinsky 1996), the set of strongly stable networks (Dutta and Mutuswami 1997; Jackson and van den Nouweland 2005), or the set of Nash networks (Bala and Goyal 2000). We also present results on the relationships between basins of attraction, the path dominance core, and the largest farsightedly consistent set (Chwe 1994). All of these results follow immediately from more general results in Page and Wooders (2005, 2008) where the notions of strong stability, pairwise stability, Nash stability, and farsighted consistency are all extended to heterogeneous directed networks.

Strongly Stable Homogeneous Networks

We begin with a formal definition of strong stability based on that of Jackson-van den Nouweland (2005),

Definition 12 (Strong Stability, Page and Wooders (2005, 2008)) Assume that G is equal to the set of all homogeneous linking networks, $P(P_2(N))$, and that the Jackson-van den Nouweland rules are in force. Network $G \in G$ is said to be strongly stable in $(G, \geq_p)$ if for all $G' \in G$ and $S \in \Gamma(N)$, $G \rightarrow_S G'$ implies that $G \not\prec_S G'$.

Thus, a network is strongly stable if whenever a coalition has the power to change the network to another network, the coalition will be deterred from doing so because the change is not preferred by the coalition. If coalitional preferences are strong, the change ‘not being preferred’ means that the change will not make *all* members of the coalition better off. If coalitional preferences are weak (i.e., based on weak preferences), the change ‘not being preferred’ means that the change will either make *no* members better off or will make some members better off and some members worse off. Note that under our definition of strong stability a network $G \in G$ that cannot be changed to another network by any coalition is strongly stable.

If (i) coalitional preferences are weak (i.e., $\{\succ_{wS}\}_{S \in \Gamma(N)}$), and (ii) coalitional effectiveness relations are determined by Jackson-van den Nouweland rules, then the definition of strong stability above is exactly that of Jackson-van den Nouweland. As it stands, our definition is closely

related to that given by Dutta-Mutuswami (Dutta and Mutuswami 1997).

We now have our main result on the path dominance core and strong stability. Denote the set of strongly stable networks by SS.

Theorem 8 (The Path Dominance Core and Strong Stability, Page and Wooders (2005, 2008)) Assume that G is equal to the set of all homogeneous linking networks, $P(P_2(N))$, and that the Jackson-van den Nouweland rules are in force.

1. If the path dominance core C of $(G, \geq_p)$ is nonempty, then SS is nonempty and $C \subseteq SS$.
2. If the dominance relation $>$ underlying $\geq_p$ is a direct dominance relation $\triangleright$, then $C = SS$ and SS is nonempty if and only if there exists a basin of attraction containing a single network.

Note that the set of strongly stable homogeneous linking networks is contained in the set of constrained Pareto efficient homogeneous linking networks. Thus, $C \subseteq SS \subseteq E$.

Pairwise Stable Networks

The following definition of pairwise stability is a translation of the Jackson-Wolinsky definition (Jackson and Wolinsky 1996).

Definition 13 (Pairwise Stability, Page and Wooders (2005, 2008)) Assume that G is equal to the set of all homogeneous linking networks, $P(P_2(N))$, and that the Jackson-Wolinsky rules are in force. Network $G \in G$ is said to be pairwise stable in $(G, \geq_p)$ if for all $\{i, i'\} \in P_2(N)$,

1. $G \rightarrow_{\{i, i'\}} G \cup \{i, i'\}$ implies that $G \not\leftarrow_{\{i, i'\}} G \cup \{i, i'\}$;
2. $G \rightarrow_{\{i\}} G \setminus \{i, i'\}$ implies that $G \not\leftarrow_{\{i\}} G \setminus \{i, i'\}$, and
3. $G \rightarrow_{\{i'\}} G \setminus \{i, i'\}$ implies that $G \not\leftarrow_{\{i'\}} G \setminus \{i, i'\}$.

Thus, a homogeneous linking network is pairwise stable if there is no incentive for any pair of players to add a link to the existing network and there is no incentive for any player who is party to a link in the existing network to dissolve or remove the link. Note that under our definition of pairwise stability a network $G \in G$ that cannot be changed to

another network by any coalition, or can only be changed by coalitions of size greater than 2, is pairwise stable.

Let PS denote the set of pairwise stable networks. It follows from the definitions of strong stability and pairwise stability that $SS \subseteq PS$. Moreover, if the full set of Jackson-Wolinsky rules are in force, then $SS = PS$. Jackson-van den Nouweland (Jackson and van den Nouweland 2005) provide two examples of the potential for strong stability to refine pairwise stability (i.e., two examples where SS is a strict subset of PS). However, under Jackson-Wolinsky rules because network changes can occur only one link at a time and because deviations by coalitions of more than two players are not possible such refinements are not possible driving SS and PS to equality.

We now have our main result on the path dominance core and pairwise stability.

Theorem 9 (The Path Dominance Core and Pairwise Stability, Page and Wooders (2005, 2008)) Assume that G is equal to the set of all homogeneous linking networks, $P(P_2(N))$, and that the Jackson-Wolinsky rules are in force.

1. If the path dominance core C of $(G, \geq_p)$ is nonempty, then PS is nonempty and $C \subseteq PS$.
2. If the dominance relation $>$ underlying $\geq_p$ is a direct dominance relation $\triangleright$, then $C = PS$ and PS is nonempty if and only if there exists a basin of attraction containing a single network. Theorem 9 can be viewed as an extension of a result due to Jackson and Watts (2002) on the existence of pairwise stable homogeneous linking networks for network formation games induced by Jackson-Wolinsky rules. In particular, Jackson and Watts (2002) show that for this particular class of Jackson-Wolinsky network formation games, if there does not exist a closed cycle of networks, then there exists a pairwise stable network. Our notion of a strategic basin of attraction containing *multiple* networks corresponds to their notion of a closed cycle of networks. Thus, stated in our terminology, Jackson and Watts show that for this class of network formation games, if there does not exist a basin of attraction

containing multiple networks, then there exists a pairwise stable network. Following our approach, by part 2 of Theorem 9 the existence of *at least one* strategic basin containing a single network is both *necessary and sufficient* for the existence of a pairwise stable network.

Nash Networks

The following definition of Nash networks is a variation on the definition from Bala and Goyal (2000).

Definition 14 (Nash Networks, Page and Wooders (2005, 2008)) Assume that G is equal to the set of all homogeneous directed networks, $P(N \times N)$, and that the Bala-Goyal rules are in force. Network $G \in G$ is said to be a Nash network in $(G, \geq_p)$ if for all $G' \in G$, $G \rightarrow_{\{i\}} G'$ implies that $G \not\prec_{\{i\}} G'$.

Thus, a homogeneous directed network is Nash if whenever an individual player has the power to change the network to another network, the player will have no incentive to do so. We shall denote by NE the set of Nash networks. Note that under our definition any network that cannot be changed to another network by a coalition of size 1 is a Nash network. Finally, note that the set of strongly stable networks SS is contained in the set of Nash networks NE.

We now have our main result on the path dominance core and Nash stability.

Theorem 10 (The Path Dominance Core and Nash Networks, Page and Wooders (2005, 2008)) Assume that G is equal to the set of all homogeneous directed networks, $P(N \times N)$, and that the Bala-Goyal rules are in force.

1. If the path dominance core C of $(G, \geq_p)$ is nonempty, then NE is nonempty and $C \subseteq \text{NE}$.
2. If the dominance relation $>$ underlying $\geq_p$ is a direct dominance relation $\triangleright$, then $C = \text{NE}$ and NE is nonempty if and only if there exists a basin of attraction containing a single network. We close this section by noting that under the Bala-Goyal rules the set of Nash networks NE is contained in the set of constrained Pareto efficient networks E. If in

addition the dominance relation is direct, then $C = \text{SS} = \text{NE} \subseteq E$.

Farsightedly Consistent Networks

Our final result summarizes the relationships between basins of attraction, the path dominance core, and the largest farsightedly consistent set.

Theorem 11 (Basins of Attraction, The Path Dominance Core, and the Largest Consistent Set, Page and Wooders (2005, 2008)) Assume that (i) G is equal to the set of homogeneous linking networks $P(P_2(N))$ and that the Jackson-Wolinsky rules or the Jackson-van den Nouweland rules are in force; or (ii) that G is equal to the set of homogeneous directed networks $P(N \times N)$ and that the Bala-Goyal rules are in force.

Given network formation game $(G, \geq_p)$, where path dominance is induced by an indirect dominance relation $\triangleright$, assume without loss of generality that $(G, \geq_p)$ has nonempty largest consistent set given by F^* and basins of attraction given by

$$\{A_1, A_2, \dots, A_m\}.$$

Then the following statements are true:

1. Each basin of attraction A_k , $k = 1, 2, \dots, m$, has a nonempty intersection with the largest consistent set F^* , that is

$$F^* \cap A_k \neq \emptyset, \quad \text{for } k = 1, 2, \dots, m.$$

2. If $(G, \geq_p)$ has a nonempty path dominance core C , then

$$C \subseteq F^*$$

Singleton Basins of Attraction

In the abstract games, $(G, \geq_p)$, that we have considered, the key condition guaranteeing nonemptiness of the path dominance core is the existence of basins of attraction containing a single network. Question: Are there classes of games

for which this is true? *In general*, if the irreflexive dominance relation $>$ inducing path dominance $\geq_p$ is transitive, then the $>$ -supernetwork $\mathbf{D} >$ is without circuits, and therefore all basins of attraction for the game $(G, \geq_p)$ contain a single network. Unfortunately, if the dominance relation is given by direct or indirect dominance, then transitivity fails to hold in general. In the next two sections we identify several classes of network formation games having singleton basins.

Network Formation Games and Potential Functions

Assume that G is equal to the set of homogeneous directed networks, $P(N \times N)$, and that the Bala-Goyal rules are in force, so that primitives are represented by supernetwork $\mathbf{G}_{bg} = \mathbf{P} \cup \mathbf{R}_{bg}$. In addition assume that player preferences over $P(N \times N)$ are specified via payoff functions $\{v_i(\cdot)\}_{i \in N}$ and that the dominance relation $>$ over $P(N \times N)$ is given by direct dominance $\triangleright$. Thus, $G' \triangleright G$ if and only if for some player $i' \in N$, $G \rightarrow \{i'\}G'$ and $v_{i'}(G') > v_{i'}(G)$.

We say that the noncooperative network formation game $(G, \geq_p)$ is a potential game if there exists a function

$$P(\cdot) : G \rightarrow \mathbb{R}$$

such that for all G and G' with $G \rightarrow \{i'\}G'$ for some player i' ,

$$v_{i'}(G') > v_{i'}(G) \text{ if and only if } P(G') > P(G).$$

It is easy to see that any noncooperative network formation game $(G, \geq_p)$ possessing a potential function (i.e., a potential game) has no circuits, and thus possesses strategic basins of attraction each consisting of a single network. Thus, we can conclude from our Theorem 8 that any noncooperative network formation game possessing a potential function has a nonempty path dominance core. In addition, we know from our Theorem 10 that in this example the path dominance core C is equal to the set of Nash networks NE .

As has been shown by Monderer and Shapley (1996), potential games are closely related to congestion games introduced by Rosenthal (1973). Page and Wooders (2007) introduce a club network formation game which is a variant of the noncooperative network formation game described above – but for a class of heterogenous directed networks – and using methods similar to those introduced by Hollard (2000), show that this game possesses a potential function. Prior papers studying potential games in the context of linking networks include Qin (1996), Slikker et al. (2000) and Slikker and van den Nouweland (2002). These papers have focused on providing the strategic underpinnings of the Myerson value (Myerson (1977) and Aumann and Myerson (1988)).

Jackson-Wolinsky Network Formation Games

Assume that G is equal to the set of all homogeneous linking networks, $P(P_2(N))$, and that the Jackson-van den Nouweland rules are in force, so that rules are represented by rules supernetwork $\mathbf{R}_{jn}$. In addition assume that player preferences over $P(P_2(N))$ are weak and therefore that coalitional preferences, $\{>_{w,S}\}_{S \in \Gamma(N)}$, are weak. Finally, assume that the dominance relation $>$ on $P(P_2(N))$ is given by direct dominance – but because coalitional preferences are weak, direct dominance is weak, denoted by $\triangleright_w$.

In the Jackson-Wolinsky network formation game coalitional preferences are specified by player payoff functions, $\{v_i(\cdot)\}_{i \in N}$, and player payoff functions are in turn specified by a network value function

$$v(\cdot) : G \rightarrow \mathbb{R}$$

together with an allocation rule, $Y(G, v) = (Y_i(G, v))_{i \in N} \in \mathbb{R}^{|N|}$ satisfying

$$\sum_{i \in N} Y_i(G, v) = v(G).$$

Thus in the Jackson-Wolinsky game, each player's payoff function $v_i(\cdot)$ is given by $Y_i(\cdot, v)$ where $v(\cdot)$ is the network value function. The basic idea here is that given network G , $v(G)$ is the total value generated by network G and $Y_i(G, v)$ is value allocated to player i . Translating Jackson-Wolinsky

into our abstract game model, if $G' \triangleright_w G$ then one of the following is true:

1. $G \rightarrow_{\{i, i'\}} G'$ where $G' = G \cup \{i, i'\}$ (a link between players i and i' is added) and $Y_i(G', v) \geq Y_i(G, v)$ and $Y_{i'}(G', v) \geq Y_{i'}(G, v)$ with strict inequality for at least one of the players;
2. $G \rightarrow_{\{i\}} G'$ where $G' = G \setminus \{i, i'\}$ (a link between players i and i' is subtracted) and $Y_i(G', v) \geq Y_i(G, v)$;
3. $G \rightarrow_{\{i\}} G'$ where $G' = G \setminus \{i, i'\}$ (a link between players i and i' is subtracted) and $Y_{i'}(G', v) > Y_{i'}(G, v)$.

Each homogeneous linking network G can be partitioned into a collection of subnetworks called components as follows. Let $H \subseteq G$ be any subnetwork of G and define

$$N_H := \{i \in N : \exists i' \in N \text{ such that } \{i, i'\} \in H\}.$$

Subnetwork $H \subseteq G$ is said to be a component if

- i and i' are in N_H , then there is a path between i and i' ; and
- if $i \in N_H$, $H \subseteq G$, and $\{i, i'\} \in G$ then $\{i, i'\} \in H$.

Let $\mathcal{C}(G)$ denote the set of all components of G .

If the allocation rule $(Y_i(G, v))_{i \in N}$, is *egalitarian* (see Jackson and Wolinsky 1996); that is, if

$$v_i(G) = Y_i(G, v) = \frac{v(G)}{|N|}$$

then any network $G^* \in \arg \max_{G \in \mathcal{G}} v(G)$ is pairwise stable and hence by part (2) of Theorem 9, $(G, \geq_p)$ has a basin of attraction containing a single network.

Alternatively, if the allocation rule $(Y_i(G, v))_{i \in N}$, is *componentwise egalitarian* (see Jackson and Wolinsky 1996); that is, if

$$\begin{aligned} v_i(G) &= Y_i(G, v) \\ &= \frac{v(H)}{|N_H|}, \text{ for } H \in \mathcal{C}(G) \text{ and } i \in N_H, \end{aligned}$$

then there exists a pairwise stable network (see Jackson 2003), and hence by part (2) of Theorem

9, $(G, \geq_p)$ has a basin of attraction containing a single network.

Finally, Jackson (2003) has shown that if the allocation rule is given by the Myerson value (see Aumann and Myerson 1988; Myerson 1977); that is if

$$v_i(G) = Y_i(G, v) = \sum_{S \subset N \setminus \{i\}} \left(v(G_{S \cup \{i\}}) - v(G_S) \right) \left(\frac{|S|!(|N| - |S| - 1)!}{|N|!} \right),$$

where

$$G_S := \{\{i, i'\} \in G : i \in S \text{ and } i' \in S\},$$

then $(G, \geq_p)$ has a basin of attraction containing a single network, and hence by part 2 of Theorem 9 $(G, \geq_p)$ has at least one pairwise stable network.

Future Directions

There are many possible directions for future research on the topic of networks and stability. Here we only mention a few that, from the perspective of this entry, seem especially promising. There are a number of potential questions to be addressed concerning the path dominance core with direct or indirect dominance. For example, what is the relationship, if any, between basins of attraction and the path dominance core and partnered (or separating) collections of coalitions, as in for example Maschler and Peleg (1967), Maschler et al. (1971), Reny and Wooders (1996) and Page and Wooders (1996)? Or what is relationship between basins of attraction and the path dominance core and the inner core, as in for example Qin (1993, 1994)?

One of the most pressing issues in our view is strategic network dynamics. Future research will address the following open question: Given the rules of network formation, the preferences of individuals, the strategic behavior of coalitions, and the trembles of nature, what network dynamics are likely to emerge and persist?

Another direction is large networks, where there are many, but still a finite number of nodes

or networks with a continuum of nodes. As in the framework of cooperative games (cf., Kovalenkov and Wooders (2001) and Wooders (1983, 2008b) for cores of transferable utility and nontransferable utility games) does it hold that some notion of the approximate path dominance core is nonempty if the numbers of players is sufficiently large? Do networks tend towards having some property analogous to the equal treatment property (as in, for example, Debreu and Scarf (1963) and Green (1972) for exchange economies, (Kovalenkov and Wooders 2001) or Wooders (2008a), for cooperative games or Gravel and Thoron (2007) for local public goods economies, or Jackson and Watts (2008) for a repeated game approach to a matching model). Then there is the problem of characterizing strategic behavior in large networks (as in, for example, Kalai (2004) or Wooders et al. (2006)). Under what conditions and to what extent might Kalai's "ex-post stability" or Wooders, Cartwright and Selter's social conformity continue to hold in strategic network formation?

Acknowledgments This paper was begun while Page and Wooders were visiting CERMSEM at the University of Paris 1 in June and October of 2007. The authors thank CERMSEM and Paris 1 for their hospitality. URLs: <http://mypage.iu.edu/~lpage>. <http://www.mymawooders.com>.

Bibliography

- Allouch N, Wooders M (2007) Price taking equilibrium in economies with multiple memberships in clubs and unbounded club sizes. *J Econ Theory*. <https://doi.org/10.1016/j.jet.2007.07.06>
- Arnold T, Wooders M (2006) Club formation with coordination. University of Warwick working paper 640
- Aumamn RJ (1964) Markets with a continuum of traders. *Econometrica* 32:39–50
- Aumamn RJ, Myerson RB (1988) Endogenous formation of links between players and coalitions: an application of the Shapley value. In: Roth A (ed) *The Shapley value*. Cambridge University Press, Cambridge, pp 175–191
- Bala V, Goyal S (2000) A noncooperative model of network formation. *Econometrica* 68:1181–1229
- Banerjee S, Konishi H, Sonmez T (2001) Core in a simple coalition formation game. *Soc Choice Welf* 18:135–158
- Belleflamme P, Bloch F (2004) Market sharing agreements and collusive networks. *Int Econ Rev* 45:387–411
- Berge C (2001) *The theory of graphs*. Dover, Mineola (reprint of the translated French edition published by Dunod, Paris, 1958)
- Bhattacharya A (2005) *Stable and efficient networks with farsighted players: the largest consistent set*. University of York, York. Typescript
- Bloch F (1995) Endogenous structures of association in oligopolies. *Rand J Econ* 26:537–556
- Bloch F (2005) Group and network formation in industrial organization: a survey. In: Demange G, Wooders M (eds) *Group formation in economics: networks, clubs, and coalitions*. Cambridge University Press, Cambridge, pp 335–353
- Bloch F, Genicot G, Ray D (2008) Informal insurance in social networks. *J Econ Theory*. <https://doi.org/10.1016/j.jet.2008.01.008>
- Blume L (1993) The statistical mechanics of strategic interaction. *Games Econ Behav* 5:387–424
- Bogomolnaia A, Jackson MO (2002) The stability of hedonic coalition structures. *Games Econ Behav* 38:201–230
- Bollobas B (1998) *Modern graph theory*. Springer, New York
- Boorman SA (1975) A combinatorial optimization model for transmission of job information through contact networks. *Bell J Econ* 6:216–249
- Bramoulle Y, Kranton R (2007a) Public goods in networks. *J Econ Theory* 135:478–494
- Bramoulle Y, Kranton R (2007b) Risk-sharing networks. *J Econ Behav Organ* 64:275–294
- Calvo-Armengol A (2004) Job contact networks. *J Econ Theory* 115:191–206
- Calvo-Armengol A, Jackson MO (2004) The effects of social networks on employment and inequality. *Am Econ Rev* 94:426–454
- Calvo-Armengol A, Jackson MO (2007) Social networks in labor markets: wage and employment dynamics and inequality. *J Econ Theory* 132:27–46
- Calvo-Armengol A, Ballester C, Zenou Y (2006) Who's who in networks. Wanted: the key player. *Econometrica* 75:1403–1418
- Casella A, Rauch J (2002) Anonymous market and group ties in international trade. *J Int Econ* 58:19–47
- Casella A, Rauch J (2003) Overcoming informational barriers in international resource allocations: prices and ties. *Econ J* 113:21–42
- Chvatal V, Lovasz L (1972) Every directed graph has a semi-kernel. In: *Hypergraph seminar, lecture notes in mathematics*, vol 411. Springer, Berlin
- Chwe M (1994) Farsighted coalitional stability. *J Econ Theory* 63:299–325
- Chwe M (2000) Communication and coordination in social networks. *Rev Econ Stud* 67:1–16
- Corominas-Bosch M (2004) Bargaining in a network of buyers and sellers. *J Econ Theory* 115:35–77
- Currarini S (2007) Group stability of hierarchies in games with spillovers. *Math Soc Sci* 54:187–202
- Currarini S, Morelli M (2000) Network formation with sequential demands. *Rev Econ Des* 5:229–249
- Debreu G, Scarf H (1963) A limit theorem on the core of an economy. *Int Econ Rev* 4:235–246

- van Deemen AMA (1991) A note on generalized stable set. *Soc Choice Welf* 8:255–260
- Demange G (1994) Intermediate preferences and stable coalition structures. *J Math Econ* 23:45–48
- Demange G (2004) On group stability and hierarchies in networks. *J Political Econ* 112:754–778
- Demange G, Henreit D (1991) Sustainable oligopolies. *J Econ Theory* 54:417–428
- Deroian F, Gannon F (2005) Quality improving alliances in differentiated oligopoly. *Int J Ind Organ* 24:629–637
- Diamantoudi E, Xue L (2003) Farsighted stability in hedonic games. *Soc Choice Welf* 21:39–61
- Durlauf S (1997) Statistical mechanics approaches to socioeconomic behavior. In: Arthur WB, Durlauf S, Lane DA (eds) *The economy as an evolving complex system II*. Addison-Wesley, Reading, pp 81–104
- Dutta B, Mutuswami S (1997) Stable networks. *J Econ Theory* 76:322–344
- Dutta B, Ghosal S, Ray D (2005) Farsighted network formation. *J Econ Theory* 122:143–164
- Even-Dar E, Kearns M, Suri S (2007) A network formation game for bipartite exchange economies. Computer and Information Science typescript, University of Pennsylvania
- Furusawa T, Konishi H (2007) Free trade networks. *J Int Econ* 72:310–335
- Galeana-Sanchez H, Xueliang L (1998) Semikernels and (k, l) -kernels in digraphs. *SIAM J Discret Math* 11:340–346
- Galeotti A, Moraga-Gonzalez JL (2007) Segmentation, advertising and prices. *Int J Ind Organ*. <https://doi.org/10.1016/j.ijindorg.2007.11.002>
- Gillies DB (1959) Solutions to general non-zero-sum games. In: Tucker AW, Luce RD (eds) *Contributions to the theory of games*, vol 4. Princeton University Press, Princeton, pp 47–85
- Goyal S (2005) Learning in networks. In: Demange G, Wooders M (eds) *Group formation in economics: networks, clubs, and coalitions*. Cambridge University Press, Cambridge, pp 122–167
- Goyal S (2007) *Connections: an introduction to the economics of networks*. Princeton University Press, Princeton
- Goyal S, Joshi S (2003) Networks of collaboration in oligopoly. *Games Econ Behav* 43:57–85
- Goyal S, Joshi S (2006) Bilateralism and free trade. *Int Econ Rev* 47:749–778
- Goyal S, Moraga-Gonzalez JL (2001) R&D networks. *Rand J Econ* 32:686–707
- Granovetter M (1973) The strength of weak ties. *Am J Sociol* 78:1360–1380
- Gravel N, Thoron S (2007) Does endogenous formation of jurisdictions lead to wealth stratification? *J Econ Theory* 132:569–583
- Green J (1972) On the inequitable nature of core allocations. *J Econ Theory* 4:132–143
- Guilbaud GT (1949) La theorie des jeux. *Econ Appl* 2:18
- Harsanyi JC (1974) An equilibrium-point interpretation of stable sets and a proposed alternative definition. *Manag Sci* 20:1472–1495
- Herings PJ-J, Mauleon A, Vannetelbosch V (2006) Farsightedly stable networks. Meteor Research Memorandum RM/06/041
- Hojman D, Szeidl A (2006) Endogenous networks, social games and evolution. *Games Econ Behav* 55:112–130
- Hollard G (2000) On the existence of a pure strategy equilibrium in group formation games. *Econ Lett* 66:283–287
- Inarra E, Kuipers J, Olaizola N (2005) Absorbing and generalized stable sets. *Soc Choice Welf* 24:433–437
- Jackson MO (2003) The stability and efficiency of economic and social networks. In: Dutta B, Jackson MO (eds) *Networks and groups: models of strategic formation*. Springer, Heidelberg, pp 99–141
- Jackson MO (2005) A survey of models of network formation: stability and efficiency. In: Demange G, Wooders M (eds) *Group formation in economics: networks, clubs, and coalitions*. Cambridge University Press, Cambridge, pp 11–57
- Jackson MO, van den Nouweland A (2005) Strongly stable networks. *Games Econ Behav* 51:420–444
- Jackson MO, Watts A (2002) The evolution of social and economic networks. *J Econ Theory* 106:265–295
- Jackson MO, Watts A (2008) Social games: matching and the play of finitely repeated games. *Games Econ Behav*. <https://doi.org/10.1016/j.geb.2008.02.004>
- Jackson MO, Wolinsky A (1996) A strategic model of social and economic networks. *J Econ Theory* 71:44–74
- Kalai E (2004) Large robust games. *Econometrica* 72:1631–1665
- Kalai E, Schmeidler D (1977) An admissible set occurring in various bargaining situations. *J Econ Theory* 14:402–411
- Kalai E, Pazner A, Schmeidler D (1976) Collective choice correspondences as admissible outcomes of social bargaining processes. *Econometrica* 44:233–240
- Kirman A (1983) Communication in markets: a suggested approach. *Econ Lett* 12:101–108
- Kirman A, Herreiner D, Weisbuch G (2000) Market organization and trading relationships. *Econ J* 110:411–436
- Konishi H, Ray D (2003) Coalition formation as a dynamic process. *J Econ Theory* 110:1–41
- Konishi H, Le Breton M, Weber S (1998) Equilibrium in a finite local public goods economy. *J Econ Theory* 79:224–244
- Kovalenkov A, Wooders M (2001) Epsilon cores of games with limited side payments: nonemptiness and equal treatment. *Games Econ Behav* 36:193–218
- Kovalenkov A, Wooders M (2003) Approximate cores of games and economies with clubs. *J Econ Theory* 110:87–120
- Kranton R, Minehart D (2000) Networks versus vertical integration. *RAND J Econ* 31:570–601
- Kranton R, Minehart D (2001) A theory of buyer-seller networks. *Am Econ Rev* 91:485–508
- Li S (1992) Far-sighted strong equilibrium and oligopoly. *Econ Lett* 40:39–44
- Li S (1993) Stability of voting games. *Soc Choice Welf* 10:51–56
- Lucas WF (1968) A game with no solution. *Bull Am Math Soc* 74:237–239

- Luo X (2001) General systems and ϕ -stable sets – a formal analysis of socioeconomic environments. *J Math Econ* 36:95–109
- Mariotti M, Xue L (2002) Farsightedness in coalition formation. Typescript, University of Aarhus
- Maschler M, Peleg B (1967) The structure of the kernel of a cooperative game. *SIAM J Appl Math* 15:569–604
- Maschler M, Peleg B, Shapley LS (1971) The kernel and bargaining set for convex games. *Int J Game Theory* 1:73–93
- Mauleon A, Varnnetelbosch V (2004) Farsightedness and cautiousness in coalition formation games with positive spillovers. *Theory Decis* 56:291–324
- Mauleon A, Sempere-Monerris J, Varnnetelbosch V (2008) Networks of knowledge among unionized firms. *Can J Econ* (to appear)
- Monderer D, Shapley LS (1996) Potential games. *Games Econ Behav* 14:124–143
- Montgomery J (1991) Social networks and labor market outcomes: toward an economic analysis. *Am Econ Rev* 81:1408–1418
- Mutuswami S, Winter E (2002) Subscription mechanisms for network formation. *J Econ Theory* 106:242–264
- Myerson RB (1977) Graphs and cooperation in games. *Math Oper Res* 2:225–229
- von Neumann J, Morgenstern O (1944) *Theory of games and economic behavior*. Princeton University Press, Princeton
- van den Nouweland A (2005) Models of network formation in cooperative games. In: Demange G, Wooders M (eds) *Group formation in economics: networks, clubs, and coalitions*. Cambridge University Press, Cambridge, pp 58–88
- Page FH Jr, Kamat S (2005) Farsighted stability in network formation. In: Demange G, Wooders M (eds) *Group formation in economics: networks, clubs, and coalitions*. Cambridge University Press, Cambridge, pp 89–121
- Page FH Jr, Wooders M (1996) The partnered core and the partnered competitive equilibrium. *Econ Lett* 52:143–152
- Page FH Jr, Wooders M (2005) Strategic basins of attraction, the farsighted core, and network formation games. FEEM Working Paper 36.05
- Page FH Jr, Wooders M (2007) Club networks with multiple memberships and noncooperative stability. Indiana University, Department of Economics typescript (paper presented at the Conference in Honor of Ehud Kalai, 16–18 Dec 2007)
- Page FH Jr, Wooders M (2008) Strategic basins of attraction, the path dominance core, and network formation games. *Games Econ Behav*. <https://doi.org/10.1016/j.geb.2008.05.003>
- Page FH Jr, Wooders M, Kamat S (2005) Networks and farsighted stability. *J Econ Theory* 120:257–269
- Qin C-Z (1993) A conjecture of Shapley and Shubik on competitive outcomes in the cores of NTU market games. *Int J Game Theory* 22:335–344
- Qin C-Z (1994) The inner core of an N-person game. *Games Econ Behav* 6:431–444
- Qin C-Z (1996) Endogenous formations of cooperation structures. *J Econ Theory* 69:218–226
- Rees A (1966) Information networks in labor markets. *Am Econ Rev* 56:218–226
- Reny PJ, Wooders M (1996) The partnered core of a game without side payments. *J Econ Theory* 70:298–311
- Richardson M (1953) Solutions of irreflexive relations. *Ann Math* 58:573–590
- Rockafellar RT (1984) *Network flows and monotropic optimization*. Wiley, New York
- Rosenthal RW (1973) A class of games possessing pure-strategy Nash equilibria. *Int J Game Theory* 2:65–67
- Roth AE (1975) A lattice fixed-point theorem with constraints. *Bull Am Math Soc* 81:136–138
- Roth AE (1977) A fixed-point approach to stability in cooperative games. In: Karamardian S (ed) *Fixed points: algorithms and applications*. Academic, New York
- Roughgarden T (2005) *Selfish routing and the price of anarchy*. MIT Press, Cambridge
- Scarf H (1967) The core of an N-person game. *Econometrica* 35:50–69
- Schwartz T (1974) Notes on the abstract theory of collective choice. Carnegie-Mellon University, School of Urban and Public Affairs typescript
- Shapley LS, Shubik M (1969) On market games. *J Econ Theory* 1:9–25
- Shenoy PP (1980) A dynamic solution concept for abstract games. *J Optim Theory Appl* 32:151–169
- Shubik M (1971) The “bridge game” economy: an example of indivisibilities. *J Political Econ* 79:909–912
- Skyrms B, Pemantle R (2000) A dynamic model of social network formation. *Proc Nat Acad Sci* 97:9340–9346
- Slikker M, van den Nouweland A (2001) *Social and economic networks in cooperative game theory*. Kluwer, Boston
- Slikker M, van den Nouweland A (2002) Network formation, costs, and potential games. In: Borm P, Peters H (eds) *Chapters in game theory*. Kluwer, Boston, pp 223–246
- Slikker M, Dutta B, van den Nouweland A, Tijs S (2000) Potential maximizers and network formation. *Math Soc Sci* 39:55–70
- Tardos E, Wexler T (2007) Network formation games and the potential function method. In: Nisan N, Roughgarden T, Tardos E, Vazirani V (eds) *Algorithmic game theory*. Cambridge University Press, Cambridge, pp 487–516
- Tesfatsion L (1997) A trade network game with endogenous partner selection. In: Amman HM, Rustem B, Whinston AB (eds) *Computational approaches to economic problems*. Kluwer, Boston, pp 249–269
- Tesfatsion L (1998) Preferential partner selection in evolutionary labor markets: a study in agent-based computational economics. In: Porto VW, Saravanan N, Waagen D, Eiben AE (eds) *Evolutionary programming VII. Proceedings of the seventh annual conference on evolutionary programming*. Springer, Berlin, pp 15–24

- Topa G (2001) Social interactions, local spillovers, and unemployment. *Rev Econ Stud* 68:261–295
- Vega-Redondo F (2007) Complex social networks. Cambridge University Press, Cambridge
- Wang P, Watts A (2006) Formation of buyer-seller trade networks in a quality differentiated product market. *Can J Econ* 39:971–1004
- Watts A (2001) A dynamic model of network formation. *Games Econ Behav* 34:331–341
- Wooders M (1983) The epsilon core of a large replica game. *J Math Econ* 11:277–300
- Wooders M (2008a) Competitive markets and market games. *Rev Econ Design* (forthcoming)
- Wooders M (2008b) Small group effectiveness, per capita boundedness and nonemptiness of approximate cores. *J Math Econ*. <https://doi.org/10.1016/j.jmateco.2007.06.006>
- Wooders M, Cartwright C, Selten R (2006) Behavioral conformity in games with many players. *Games Econ Behav* 57:347–360
- Xue L (1998) Coalitional stability under perfect foresight. *Econ Theory* 11:603–627
- Xue L (2000) Negotiation-proof Nash equilibrium. *Int J Game Theory* 29:339–357
- Zissimos B (2005) Why are free trade agreements regional? FEEM Working Paper 67-07

Game Theory and Strategic Complexity

Kalyan Chatterjee¹ and Hamid Sabourian²

¹Department of Economics, The Pennsylvania State University, University Park, USA

²Faculty of Economics, University of Cambridge, Cambridge, UK

Article Outline

Glossary

Definition

Introduction

Games, Automata, and Equilibrium Concepts

Complexity Considerations in Repeated Games

Complexity and Bargaining

Complexity, Market Games, and the Competitive Equilibrium

Discussion and Future Directions

Bibliography

Glossary

Automata A formal definition of a strategy that captures its complexity.

Continuation Game A description of how the play will proceed in a dynamic game once some part of the game has already occurred.

Equilibrium A solution concept for games in which each player optimizes given his correct prediction of others' behavior.

Equilibrium Path The outcome in terms of the play of the game if every player uses his equilibrium strategy.

Game Theory A formal model of interaction, usually in human behavior.

Repeated Games A series of identical interactions of this kind.

Strategic Complexity A measure of how complex a strategy is to implement.

Strategy A complete specification of how a player will play the game.

Definition

The subject of this entry is at the intersection of economics and computer science and deals with the use of measures of complexity obtained from the study of finite automata to help select among multiple equilibria and other outcomes appearing in game-theoretic models of bargaining, markets, and repeated interactions. The importance of the topic lies in the ability of concepts that employ bounds on available resources to generate more refined predictions of individual behavior in markets.

Introduction

This entry is concerned with the concept of strategic complexity and its use in game theory. There are many different meanings associated with the word “complexity,” as the variety of topics discussed in this volume makes clear. In this entry, we shall adopt a somewhat narrow view, confining ourselves to notions that measure, in some way, constraints on the ability of economic agents to behave with full rationality in their interactions with other agents in dynamic environments. This will be made more precise a little later. (A more general discussion is available in Rubinstein (1998)).

Why is it important to study the effect of such constraints on economic decision-making? The first reason could be to increase the realism of the assumptions of economic models; it is evident from introspection and from observing others that we do not have infinite memory and cannot condition our future actions on the entire corpus of what we once knew or, for that matter, unlimited computational power. However, only considering the assumptions of a model would not be considered enough if the increased realism were not to expand our ability to explain or to predict. The second reason therefore is that studying the effects of complexity on human decision-making might help us either to make our predictions more

precise (by selecting among equilibria) or to generate explanations for behavior that is frequently observed, but incompatible with equilibrium in models that have stronger assumptions about the abilities of agents.

A strategy in a game is an entire plan of how to play the game at every possible history/contingency/eventuality at which the player has to make a move. The particular aspect of complexity that we shall focus on is on the complexity of strategy as a function of the history. One representation of the players' strategies in games is often in terms of (finite) automata. The finiteness need not always be assumed; it can be derived. The ideas of complexity, though often most conveniently represented this way, can also be discussed without referring to finite automata at all but purely in how a strategy depends on the past history of the game.

The number of states in the automaton can be used as a measure of complexity. This may be a natural measure of complexity in a stationary repetitive environment such as repeated games. We shall discuss this measure of complexity as well as other aspects of the complexity of a strategy that are particularly relevant in non-stationary frameworks.

Note the players are not themselves considered automata in this entry, and in the literature it surveys. Also, we do not place restrictions on the ability of players to compute strategies (see Papadimitriou 1992), only on the strategies that they can implement. The entry is also not intended as a comprehensive survey of the literature on complexity of implementation in games. The main focus of the entry is inevitably on the works that we have been personally associated with.

The remaining part of this entry is organized as follows: In the next section, we discuss strategies in a game, their representation as finite automata, and the basic equilibrium concepts to be used in the entry. Section "[Complexity Considerations in Repeated Games](#)" will consider the use of complexity notions in repeated games. Section "[Complexity and Bargaining](#)" will focus on extensive form bargaining and the effect of complexity considerations in selecting equilibria. Section "[Complexity, Market Games and the](#)

[Competitive Equilibrium](#)" will extend the analysis of bargaining to markets in which several agents bargain and consider the recent literature that justifies competitive outcomes in market environments by appealing to the aversion of agents to complexity. Section "[Discussion and Future Directions](#)" concludes with some thoughts on future research. This entry draws on an earlier survey paper (Chatterjee 2002) for some of the material in sections "[Games, Automata and Equilibrium Concepts](#)," "[Complexity Considerations in Repeated Games](#)," and "[Complexity and Bargaining](#)."

Games, Automata, and Equilibrium Concepts

As mentioned in the introduction, this entry will be concerned with *dynamic games*. Though the theory of games has diffused from economics and mathematics to several other fields in the last few decades, we include an introduction to the basic concepts to keep this entry as self-contained as possible. A game is a formal model of interaction between individual agents. The basic components of a game are the following: (i) Players or agents, whose choices will, in general, have consequences for each other. We assume a finite set of players, denoted by N . We shall also use N sometimes to represent the cardinality of this set. (ii) A specification of the "rules of the game" or the structure of interaction, described by the sequence of possible events in the game, the order in which the players move, what they can choose at each move, and what they know about previous moves. This is usually modeled as a tree and is called the "extensive form" of the game (and will not be formalized here, though the formalization is standard and found in all the texts on the subject). (iii) Payoffs for each player associated with every path through the tree from the root. It is easier to describe this as a finite tree and ascribe payoffs to the end nodes z . Let $u_i(z)$ be the real-valued payoff to Player i associated with end node z . The payoffs are usually assumed to satisfy conditions that are sufficient to guarantee that the utility of a probability distribution on a subset

of the set of end nodes is the expectation of the utility of the individual end nodes. However, different strands of work on bounded rationality dispense with this assumption. The description above presupposes a tree of finite depth, while many of the applications deal with infinite horizon games. However, the definitions are easily modified by associating payoffs with a *play* of the game and defining a node as a set of plays. We shall not pursue this further here.

In the standard model of a game, players are assumed to have all orders of knowledge about the preceding description. Work on bounded rationality also has considered relaxing this assumption.

A strategy is a complete plan of action for playing a game, describing the course of action to be adopted in every possible contingency (or every information set of the player concerned). The plan has to be detailed enough so that it can be played by an agent, even if the principal is not himself or herself in town, and the agent could well be a computer, which is programmed to follow the strategy. Without any loss of generality, a strategy can be represented by an automaton (see below for illustration and Osborne and Rubinstein (1994) for a formal treatment in the context of repeated games). Often such a machine description is more convenient in terms of accounting for a complexity of a machine. For example, the works that are based on the use of finite automata or Turing machines to represent *strategies* for playing a game impose a natural bound on the set of allowable strategies.

For the types of problem that we shall consider here, it is best to think of a *multistage game with observable actions*, to use the terminology of Fudenberg and Tirole (1991). The game has some temporal structure; let us call each unit of time a period or a stage. In each period, the players choose actions simultaneously and independently. (The actions could include the dummy action.) All the actions taken in a stage are observed, and the players then choose actions again. An example is a repeated normal-form game, such as the famous Prisoners' Dilemma being repeated infinitely or finitely often. In each stage, players choose whether to *cooperate* or *defect*. The choices are revealed, payoffs received and the choices

repeated again, and so on. (The reader will recall that in the Prisoners' Dilemma played once, *Defect* is better than *Cooperate* for each player, no matter what the other player does, but both players choosing *Defect* is strictly worse for each than both choosing *Cooperate*.) What a strategy for a given player would do would be to specify the choice in a given stage as a function of the history of the game up to that stage (for every stage). A finite automaton represents a particular strategy in the following way: It partitions all possible histories in the game (at which the player concerned has to move) using a finite number of elements. Each of these elements is a *state* of the machine. Given a state, the automaton prescribes an action (e.g., Cooperate after all histories in which the other party has cooperated). It also specifies how the state of the machine will change as a result of the action taken by the other player. The state-to-action mapping is called the output mapping, and the rule that prescribes the state in the next period as a function of today's state and the action of one's opponent in this period is called the transition mapping. The automaton also needs to prescribe what to do in the first stage, when there is no history of past actions to rely on. Thus, for example, the famous "tit-for-tat" strategy in the repeated Prisoners' Dilemma can be represented by the following automaton:

1. Play Cooperate in the first stage. The initial state is denoted as q^1 , and in this state, the action prescribed is cooperate.
2. As long as the other player cooperates, stay in state q^1 .
3. If the other player defects in a state, go to state q^2 . The action specified in q^2 is *Defect*.
4. Stay in q^2 as long as the other player defects. If the other player cooperates in a stage, go to q^1 .

Denoting the output mapping by $\lambda(\cdot)$, we get $\lambda(q^1) = C$ and $\lambda(q^2) = D$. The transition mapping $\mu(\cdot, \cdot)$ is as follows: $\mu(q^1, C) = q^1$, $\mu(q^1, D) = q^2$, $\mu(q^2, C) = q^1$, and $\mu(q^2, D) = q^2$. Here, of course, C and D denote cooperate and defect, respectively. The machine described above has two states and is an instance of a Moore machine in computer science terminology.

The use of a Moore machine to represent a strategy rules out strategies in which histories are arbitrarily finitely partitioned or arbitrarily complex. In fact, the number of states in the machine is a popular measure of the complexity of the machine and the strategy it represents.

Another kind of finite automaton used in the literature is a Mealy machine. The main difference between this and the Moore machine is that now the output is a function both of the state and of an input, unlike the Moore machine where it is only a function of the state. One can always transform a Mealy machine to a Moore machine by making transitions depend on the input and having state transitions after every input. The Mealy machine representation is more convenient for the extensive form game we shall consider in section “Complexity and Bargaining.” We shall briefly address why in that section.

The aim of using the machine framework to describe strategies is to take into account explicitly the cost of complexity of strategies. There is the belief, for instance, that short-term memory (see Miller 1956) is capable of keeping seven things in mind at any given time, and if five of them are occupied by how to play the Prisoners’ Dilemma, there might be less leftover for other important activities.

The standard equilibrium concept in game theory is the concept of Nash equilibrium. This requires each player to choose a best strategy (in terms of payoff) given his or her conjectures about other players’ strategies, and, of course, in equilibrium the conjectures must be correct. Thus, a Nash equilibrium is a profile of strategies, one for each player, such that every player is choosing the best response strategy given the Nash equilibrium strategies of the other players. In dynamic games Nash equilibrium strategies may not be credible (sequentially rational). In multistage games, to ensure credibility, the concept of Nash equilibrium is refined by requiring the strategy of each player to be a best response to the strategies of the others at every well-defined history (subgame) within the game. This notion of equilibrium was introduced by Selten (1965) and is called subgame perfect equilibrium. The difference between this concept and that of Nash,

which it refines, is that players must specify strategies that are best responses to each other even at nodes in the game tree that would never be reached if the prescribed equilibrium were being played. The Nash concept does not require this. The notion of histories *off the equilibrium path* therefore refers to those that do not occur if every player follows his or her equilibrium strategy. Another useful concept to mention here is that of *payoff in the continuation game*. This refers to the expected payoff from the prescribed strategies in the part of the game remaining to be played after some moves have already taken place. The restriction of the prescribed strategies to the continuation game is referred to here as *continuation strategies*.

Rubinstein (1986), Abreu and Rubinstein (1988), and others have modified the standard equilibrium concepts to account for complexity costs. This approach is somewhat different from that adopted, for example, by Neyman (1985), who restricted strategies to those of bounded complexity. We shall next present the Abreu-Rubinstein definition of *Nash equilibrium with complexity* (often referred to as NEC in the rest of the entry).

The basic idea is a very simple extension of Nash equilibrium. Complexity enters the utility function lexicographically. A player first calculates his or her best response to the conjectured strategies of the other players. If there are alternative best responses, the player chooses the less complex one. Thus, a Nash equilibrium with complexity has two aspects. First, the strategies chosen by any player must be the best response given his or her conjectures about other players’ strategies, and, of course, in equilibrium the conjectures must be correct. Second, there must not exist an alternative strategy for a player such that his or her payoff is the same as in the candidate equilibrium strategy, given what other players do, but the alternative strategy is less complex.

In Abreu and Rubinstein (1988), the measure of complexity is the number of states in the Moore machine that represents the strategy. The second part of their equilibrium definition restricts the extent to which punishments can be used off the equilibrium path. For example, there is a famous

strategy that, if used by all players, gives cooperation in the infinitely repeated Prisoners' Dilemma (for sufficiently high discount factors), namely, the "grim" strategy. This strategy can be described by the following machine: Start with *Cooperate*. Play *Cooperate* as long as the other players all cooperate. If in the last period any player has used *Defect*, then switch to playing *Defect* forever (i.e., never play *Cooperate* again, no matter what the other players do in succeeding periods). This strategy profile (each player uses the grim strategy) gives an outcome path consisting solely of players cooperating. No one defects because from then until the end of time all the players will be punishing one another.

However, this strategy profile is not a Nash equilibrium with complexity; the grim strategy is a two-state machine in which one state (the one in which a player chooses *Defect*) is never used given that everyone else cooperates on the equilibrium path. Some player can do better, even if lexicographically, by switching to a one-state machine in which he or she cooperates no matter what. Thus, even the weak lexicographic requirement has some bite.

Note that the complexity restriction we are considering is on the complexity of *implementation*, not the complexity of *computation*. We know that even a Turing machine, which has potentially infinite memory, might be unable to calculate best responses to all possible strategy profiles of other players in the game (see Anderlini 1990; Binmore 1987).

To return to the question of defining equilibrium in the machine game, the Abreu-Rubinstein approach is described by them as "buying" states in the machine at the beginning of the game. The complexity cost is therefore a fixed cost per state used. Some recent papers have taken the fixed-cost approach further by requiring NEC strategies to be credible. The idea is that players pay an initial fixed cost for the complexity (the notion of complexity in some of these papers differs from counting the states approach) of his/her strategy and then the game is played with strategies being optimal at every contingency as in standard game theory. Chatterjee and Sabourian (2000a, b) model this by considering Nash equilibrium with

complexity costs in (bargaining) games in which machines/strategies can make errors/trembles in output/action. The introduction of errors ensures that the equilibrium strategies are optimal after every history. As the error goes to zero, we are left with subgame perfect equilibria of the underlying game. Chatterjee and Sabourian (2000a), Sabourian (2003), Gale and Sabourian (2005), and Lee and Sabourian (2007) take a more direct method of introducing credibility into the equilibrium concept with complexity costs by restricting NEC strategies to be subgame perfect equilibrium in the underlying game with no complexity costs. We refer to such an equilibria by perfect equilibrium with complexity costs (PEC).

In contrast to the fixed-cost interpretation of complexity cost, Rubinstein in his 1986 paper considers a different approach, namely, the choice of "renting" states in the machine for every period the game is played. Formally, the Rubinstein notion of *semi-perfect equilibrium* requires the strategy chosen to have the minimal number of states necessary to play the game at every node on the (candidate) equilibrium outcome path. A state could therefore be dropped if it is not going to be used on the candidate equilibrium path after some period. Thus, to be in the equilibrium machine, it is not sufficient that a state be used on the path; it has to be used in every possible future. Rubinstein called this notion of equilibrium *semi-perfect*, because the complexity of a strategy could be changed in one direction (it could be decreased) after every period. If states could be added as well as deleted every period, we would have yet another definition of equilibrium with complexity, *machine subgame perfect equilibrium*. (See Neme and Quintas 1995.) In contrast, both the NEC and PEC concepts we use here entail a *single choice* of automaton or strategy by players at the beginning of the game.

In all these models, complexity analysis has been facilitated by considering the "machine games." Each player chooses among machines, and the complexity of a machine is taken to be the number of states of the machine. In fact, the counting-the-number-of-states measure of complexity has an equivalent measure stated in terms of the underlying strategies that the machine could

implement. Kalai and Stanford (1988) define complexity of a strategy by the number of *continuation strategies* that the strategy induces at different periods/histories of the game and establishes that such a measure is equal to the number of the states of the smallest machine that implements the strategy. Thus, one could equivalently describe any result either in terms of underlying strategies and the cardinality of the set of continuation strategies that they induce or in terms of machines and the number of states in them. The same applies to other measures of complexity discussed in this entry; they can be defined either in terms of the machine specification or in terms of the underlying strategy. In the rest of this entry, to simplify the exposition, we shall at times go from one exposition to the other without further explanation.

With this preamble on the concepts of equilibrium used in this literature, we turn to a discussion of a specific game in the next section, the infinitely repeated Prisoners' Dilemma. We will discuss mainly the approach of Abreu and Rubinstein in this section but contrast it with the literature following from Neyman. We also note that the suggestion for using finite automata in games of this kind came originally from Aumann (1981).

Complexity Considerations in Repeated Games

Endogenous Complexity

In this subsection, we shall first concentrate on the Prisoners' Dilemma and discuss the work of Abreu and Rubinstein, which was introduced briefly in the last section. For concreteness, consider the following Prisoners' Dilemma payoffs:

	C_2	D_2
C_1	3, 3	-1, 4
D_1	4, -1	0, 0

This is the "stage game"; each of the two players chooses an action in each stage; their actions are revealed at the end of the stage, and then the next stage begins. The game is repeated

infinitely often, and future payoffs are discounted with a common discount factor δ .

The solution concept to be used was introduced in the last section, NEC or Nash equilibrium with complexity. Note that here complexity is *endogenous*. A player has a preference for less complex strategies. This preference comes into play lexicographically, that is, for any strategies or machines that give the same payoff against the opponent's equilibrium strategy, a player will choose the one with lowest complexity. Thus, the cost of complexity is infinitesimal. One could also consider positive but small costs of more complex strategies, but results will then depend on how large the cost of additional complexity is compared to the additional payoff obtained with a more complex strategy.

We saw in the last section that the "grim trigger" strategy, which is a two-state automaton, is not an NEC. The reason is that if Player 2 uses such a strategy, Player 1 can be better off by deviating to a one-state strategy in which she always cooperates. (This will give the same payoff with a less complex strategy.) One-state strategies where both players cooperate clearly do not constitute NEC (deviating and choosing a one-state machine that always plays D is strictly better for a player). However, if both players use a one-state machine that always generates an action of D , this is an NEC.

The question obviously arises if the cooperative outcome in each stage can be sustained as an NEC, and the preceding discussion makes clear that the answer is no. Punishments have to be used on the equilibrium path, but we can get arbitrarily close to the cooperative outcome for a high enough discount factor. For example, consider the following two-state machine:

$$\begin{aligned} Q &= \{q^1, q^2\}; \lambda(q^1) = D, \lambda(q^2) = C, \\ \mu(q^1, C) &= q^1, \mu(q^1, D) = q^2, \mu(q^2, D) = q^1, \\ \mu(q^2, C) &= q^2. \end{aligned}$$

Here both players play the same strategy, which starts out playing D . If both players do as they are supposed to, each plays C in the next period and thereafter, so the sequence of actions

is (D,D) , (C,C) , (C,C) If either player plays C in the first period, the other player keeps playing D in the next period. The transition rule prescribes that if one plays C and one's opponent plays D , one goes back to playing D , so the sequence with the deviation will be (D,C) , (D,D) , (C,C) , (C,C)

Suppose both players use this machine. First, we check it is a Nash equilibrium in payoffs. We only need to check what happens when a player plays C . If Player 2 deviates and plays D , she will get an immediate payoff of 4 followed by payoffs of 0, 3, 3. . . if she thereafter sticks to her strategy for a total payoff of $4 + \delta^2(3/(1 - \delta))$ as opposed to $3/(1 - \delta)$ if she had not deviated. The net gain from deviation is $1 - 3\delta$, which is negative for $\delta > \frac{1}{3}$. One can check that more complicated deviations are also worse. The second part of the definition needs to be checked as well, so we need to ensure that a player cannot do as well in terms of payoff by moving to a less complex strategy, namely, a one-state machine. A one-state machine that always plays C will get the worst possible payoff, since the other machine will keep playing D against it. A one-state machine that plays D will get a payoff of 4 in periods 2, 4, 6. . . or a total payoff of $4\delta/(1 - \delta^2)$ as against $3\delta/(1 - \delta)$. The second is strictly greater for $\delta > \frac{1}{3}$.

This machine gives a payoff close to 3 per stage for δ close to 1. As $\delta \rightarrow 1$, the payoff of each player goes to 3, the cooperative outcome.

The paper by Abreu and Rubinstein obtains a basic result on the characterization of payoffs obtained as NEC in the infinitely repeated Prisoners' Dilemma. We recall that the "Folk Theorem" for repeated games tells us that all outcome paths that give a payoff per stage strictly greater for each player than the minmax payoff for that player in the stage game can be sustained by Nash equilibrium strategies. Using endogenous complexity, one can obtain a refinement; now only payoffs on a so-called cross are sustainable as NEC. This result is obtained from two observations. First, in any NEC of a two-player game, the number of states in the players' machines must be equal. This follows from the following intuitive reasoning (we refer readers to the original paper for the proofs). Suppose we fix the machine used

by one of the players (say Player 1), so that to the other player it becomes part of the "environment." For Player 2 to calculate the best response or an optimal strategy to Player 1's given machine, it is clearly not necessary to partition past histories more finely than the other player has done in obtaining her strategy; therefore, the number of states in Player 2's machine need not (and therefore will not, if there are complexity costs) exceed the number in Player 1's machine in equilibrium. The same holds true in the other direction, so the number of states must be equal. (This does not hold for more than two players.) Another way of interpreting this result is that it restates the result from Markov decision processes on the existence of an optimal "stationary" policy (i.e., depending only on the states of the environment, which are here the same as the states of the other player's machine). See also Piccione (1992).

Thus, there is a one-to-one correspondence between the states of the two machines. (Since the number of states is finite and the game is infinitely repeated, the machine must visit at least one of the states infinitely often for each player.) One can strengthen this further to establish a one-to-one correspondence between *actions*. Suppose Player 1's machine has $a_t^1 = a_s^1$, where these denote the actions taken at two distinct periods and states by Player 1, with $a_t^2 \neq a_s^2$ for Player 2. Since the states in t and s are distinct for Player 1 and the actions taken are the same, the transitions must be different following the two distinct states. But then Player 1 does not need two distinct states; he can drop one and condition the transition after, say, s on the different action used by Player 2. (Recall the transition is a function of the state and the opponent's action.) But then Player 1 would be able to obtain the same payoff with a less complex machine; so the original one could not have been an NEC machine.

Therefore, the actions played must be some combination of (C,C) and (D,D) (the correspondence is between the two C s and the two D s) or some combination of (C,D) and (D,C) . (By combination, we mean combination over time. For example, (C,C) is played, say, 10 times for every 3 plays of (D,D) . In the payoff space,

sustainable payoffs are either on the line joining (3,3) and (0,0) or on the line joining the payoffs on the other diagonal, hence the evocative name chosen to describe the result – the cross of the two diagonals.

While this is certainly a selection of equilibrium outcomes, it does not go as far as we would wish. We would hope that some equilibrium selection argument might deliver us the cooperative outcome (3,3) uniquely (even in the limit as $\delta \rightarrow 1$), instead of the actual result obtained. There is work that does this, but it uses evolutionary arguments for equilibrium selection (see Binmore and Samuelson (1992)). An alternative learning argument for equilibrium selection is used by Maenner (2008). In his model, a player tries to infer what machine is being used by his opponent and chooses the simplest automaton that is consistent with the observed pattern of play as his model of his opponent. A player then chooses the best response to this inference. It turns out complexity is not sufficient to pin down an inference, and one must use optimistic or pessimistic rules to select among the simplest inferences. One of these gives only (D,D) repeated, while the other reproduces the Abreu-Rubinstein NEC results. Piccione and Rubinstein (1993) show that the NEC profile of 2-player repeated *extensive form games* is unique if the stage game is one of perfect information. This unique equilibrium involves all players playing their one-shot myopic noncooperative actions at every stage. This is a strong selection result and involves stage game strategies not being observable (only the path of play is) as well as the result on the equilibrium numbers of states being equal in the two players' machines.

In repeated games with more than two players or with more than two actions at each stage, the multiplicity problem may be more acute than just not being able to select uniquely a "cooperative outcome." In some such games, complexity by itself may not have any bite, and the Folk theorem may survive even when the players care for the complexity of their strategies. (See Bloise 1998 who shows robust examples of two-player repeated games with three actions at each stage such that every individually rational payoff can be sustained as an NEC if players are sufficiently patient.)

Exogenous Complexity

We now consider the different approach taken by Neyman (1985, 1997), Ben Porath (1986, 1993), Zemel (1989), and others. We shall confine ourselves to the papers by Neyman and Zemel on the Prisoners' Dilemma, without discussing the more general results these authors and others have obtained.

Neyman's approach treats complexity as exogenous complexity. Let Player i be restricted to use strategies/automata with the number of states not to exceed m_i . He also considers finitely repeated games, unlike the infinitely repeated games we have discussed up to this point. With the stage game being the Prisoners' Dilemma and the number of repetitions being T (for convenience, this includes the first time the game is played, we can write the game being considered as $G^T(m_1, m_2)$). Note that without the complexity restrictions, the finitely repeated Prisoners' Dilemma has a unique Nash equilibrium outcome path (and a unique subgame perfect equilibrium)-(D,D) in all stages. Thus, sustaining cooperation in this setting is obtaining non-equilibrium behavior, though one that is frequently observed in real life. This approach therefore is an example of bounded rationality being used to explain observed behavior that is not predicted in equilibrium.

If the complexity restrictions are severe, it turns out that (C,C) in each period is an equilibrium. For this, we need $2 \leq m_1, m_2 \leq T - 1$. To see this consider the grim trigger strategy mentioned earlier – representable as a two-state automaton – and let $T = 3$. Here $\lambda(q^1) = C$; $\lambda(q^2) = D$; $\mu(q^1, C) = q^1$; $\mu(q^1, D) = q^2$; $\mu(q^2, C \text{ or } D) = q^2$. If each player uses this strategy, (C,C) will be observed. Such a pair of strategies is clearly not a Nash equilibrium-given Player 1's strategy; Player 2 can do better by playing D in stage 3. But if Player 2 defects in the second stage, by choosing a two-state machine where $\mu(q^1, C) = D$, he will gain 1 in the second stage and lose 3 in the third stage as compared to the machine listed above, so he is worse off. But defecting in stage 3 requires an automaton with three states – two states in which C is played and one in which D is played. The transitions in state q^1 will be similar, but, if q^2 is the second

cooperative state, the transition from q^2 to the defect state will take place no matter whether the other player plays C or D . However, automata with three states violate the constraint that the number of states be no more than 2, so the profitable deviation is out of reach.

While this is easy to see, it is not clear what happens when the complexity is high. Neyman shows the following result: *For any integer k , there exists a T_0 , such that for $T \geq T_0$ and $T^{\{(1/k)\}} \leq m_1, m_2 \leq T^{\{k\}}$ there is a mixed strategy equilibrium of $G^T(m_1, m_2)$ in which the expected average payoff to each player is at least $3 - \frac{1}{k}$.*

The basic idea is that rather than playing (C, C) at each stage, players are required to play a complex sequence of C and D , and keeping track of this sequence uses up a sufficient number of states in the automaton so that profitable deviations again hit the constraint on the number of states. But since D cannot be avoided on the equilibrium path, only something close to (C, C) each period can be obtained rather than (C, C) all the time.

Zemel's paper adds a clever little twist to this argument by introducing communication. In his game, there are two actions each player chooses at each stage, either C or D as before and a message to be communicated. The message does not directly affect payoffs as the choice of C or D does. The communication requirements are now made sufficiently stringent, and deviation from them is considered a deviation, so that once again the states "left over" to count up to N are inadequate in number and (C, C) can once again be played in each stage/period. This is an interesting explanation of the rigid "scripts" that many have observed to be followed, for example, in negotiations.

Neyman (1997) surveys his own work and that of Ben Porath (1986, 1993). He also generalizes his earlier work on the finitely repeated Prisoners' Dilemma to show how small the complexity bounds would have to be in order to obtain outcomes outside the set of (unconstrained) equilibrium payoffs in the finitely repeated, normal-form game (just as (C, C) is not part of an unconstrained equilibrium outcome path in the Prisoners' Dilemma). Essentially, if the complexity permitted grows exponentially or faster with the number of repetitions, the equilibrium payoff sets of the

constrained and the unconstrained games will coincide. For sub-exponential growth, a version of the Folk theorem is proved for two-person games. The first result says:

For every game G in strategic form and with m_i being the bound on the complexity of i 's strategy and T the number of times the game G is played, there exists a constant c such that if $m_i \geq \exp(ct)$, then $E(G^T) = E(G^T(m_1, m_2))$ where $E(\cdot)$ is the set of equilibrium payoffs in the game concerned.

The second result, which generalizes the Prisoners' Dilemma result already stated, considers a sequence of triples $(m_1(n), m_2(n), T(n))$ for a two-player strategic form game, with $m_2 \geq m_1$, and shows that the $\liminf$ of the set of equilibrium payoffs of the automata game as $n \rightarrow \infty$ includes essentially the strictly individually rational payoffs of the stage game if $m_1(n) \rightarrow \infty$ and $(\log m_1(n))/T(n) \rightarrow 0$ as $n \rightarrow \infty$. Thus, a version of the Folk theorem holds provided the complexity of the players' machines does not grow too fast with the number of repetitions.

Complexity and Bargaining

Complexity and the Unanimity Game

The well-known alternating offers bargaining model of Rubinstein has two players alternating in making proposals and responding to proposals. Each period or unit of time consists of one proposal and one response. If the response is "reject," the player who rejects makes the next proposal but in the following period. Since there is discounting with discount factor δ per period, a rejection has a cost. The unanimity game we consider is a multi-person generalization of this bargaining game, with n players arranged in a fixed order, say $1, 2, 3, \dots, n$. Player 1 makes a proposal on how to divide a pie of size unity among the n people; players $2, 3, \dots, n$ respond sequentially, either accepting or rejecting. If everyone accepts, the game ends. If someone rejects, Player 2 now gets to make a proposal but in the next period. The responses to Player 2's proposal are made sequentially by Players $3, 4, 5, \dots, n, 1$. If Player i gets a share x_i in an eventual agreement at time t , his payoff is $\delta^{t-1}x_i$.

Avner Shaked had shown in 1986 that the unanimity game had the disturbing feature that all individually rational (i.e., nonnegative payoffs for each player) outcomes could be supported as subgame perfect equilibria. Thus, the sharp result of Rubinstein (1982), who found a unique subgame perfect equilibrium in the two-player, stood in complete contrast with the multiplicity of subgame perfect equilibria in the multiplayer game.

Shaked's proof had involved complex changes in expectations of the players if a deviation from the candidate equilibrium were to be observed. For example, in the three-player game with common discount factor δ , the three extreme points $(1,0,0)$, $(0,1,0)$, and $(0,0,1)$ sustain one another in the following way. Suppose Player 1 is to propose $(0,1,0)$, which is not a very sensible offer for him or her to propose, since it gives everything to the second player. If Player 1 deviates and proposes, say, $((1 - \delta)/2, \delta, (1 - \delta)/2)$ then it might be reasoned that Player 2 would have no incentive to reject because in any case he or she can't get more than 1 in the following period, and Player 3 would surely prefer a positive payoff to 0. However, there is a counterargument. In the subgame following Player 1's deviation, Player 3's expectations have been raised so that he (and everyone else, including Player 1) now expect the outcome to be $(0,0,1)$, instead of the earlier expected outcome. For sufficiently high discount factor, Player 3 would reject Player 1's insufficiently generous offer. Thus, Player 1 would have no incentive to deviate. Player 1 is thus in a bind; if he offers Player 2 less than δ and offers Player 3 more in the deviation, the expectation that the outcome next period will be $(0,1,0)$ remains unchanged, so now Player 2 rejects his offer. So no deviation is profitable, because each deviation generates an expectation of future outcomes, an expectation that is confirmed in equilibrium. (This is what equilibrium means.) Summarizing, $(0,1,0)$ is sustained as follows: Player 1 offers $(0,1,0)$, and Player 2 accepts any offer of at least 1 and Player 3 any offer of at least 0. If one of them rejects Player 1's offer, the next player in order offers $(0,1,0)$ and the others accept. If any proposer, say Player 1, deviates from the offer $(0,1,0)$ to (x_1, x_2, x_3) , the player with the lower of $\{x_2, x_3\}$

rejects. Suppose it is Player i who rejects. In the following period, the offer made gives 1 to Player i and 0 to the others, and this is accepted.

Various attempts were made to get around the continuum of equilibria problem in bargaining games with more than two players; most of them involved changing the game. (See Chatterjee and Sabourian 2000a, b for a discussion of this literature.) An alternative to changing the game might be to introduce a cost for this additional complexity, in the belief that players who value simplicity will end up choosing simple, that is, history-independent strategies. This seems to be a promising approach because it is clear from Shaked's construction that the large number of equilibria results from the players choosing strategies that are history dependent. In fact, if the strategies are restricted to those that are history independent (also referred to as *stationary* or *Markov*), then it can be shown (see Herrero 1985) that the subgame perfect equilibrium is unique and induces equal division of the pie as $\delta \rightarrow 1$.

The two papers (Chatterjee and Sabourian 2000a, b) in fact seek to address the issue of complex strategies with players having a preference for simplicity, just as in Abreu and Rubinstein. However, now we have a game of more than two players and a single extensive form game rather than a repeated game as in Abreu-Rubinstein. It was natural that the framework had to be broadened somewhat to take this into account.

For each of n players playing the unanimity game, we define a machine or an implementation of the strategy as follows.

A *stage* of the game is defined to be n periods, such that if a stage were to be completed, each player would play each role at most once. A *role* could be as proposer or $n-1$ th responder or $n-2$ th responder . . . up to the first responder (the last role would occur in the period before the player concerned had to make another proposal). An *outcome* of a stage is defined as a sequence of offers and responses, for example, $e = (x, A, A, R; y, R; z, A, R; b, A, A, A)$ in a four-player game where the (x, y, z, b) are proposals made in the four periods and (A, R) refer to accept and reject, respectively. From the point of view of the first player to

propose (for convenience, let's call him Player 1), he makes an offer x , which is accepted by Players 2 and 3 but rejected by Player 4. Now it is Player 2's turn to offer, but this offer, y , is rejected by the first responder Player 3. Player 1 gets to play as second responder in the next period, where he rejects Player 3's proposal. In the last period of this stage, a proposal b is made by Player 4 and everyone accepts (including Player 1 as first responder). Any partial history within a stage is denoted by s . For example, when Player 2 makes an offer, he does so after a partial history $s = (x, A, A, R)$. Let the set of possible outcomes of a stage be denoted by E and the set of possible partial histories by S . Let Q_i denote the set of states used in the i th player's machine M_i . The output mapping is given by $\lambda_i: S \times Q_i \rightarrow \Lambda$, where Λ is the set of possible actions (i.e., the set of possible proposals, plus accept or reject). The transition between states now takes place at the end of each stage, so the transition mapping is given as $\mu_i: E \times Q_i \rightarrow Q_i$. As before, in the Abreu-Rubinstein setup, there is an initial state $q_{\text{initial},i}$ specified for each player. There is also a termination state F , which is supposed to indicate agreement. Once in the termination state, players will play the null action and make transitions to this state.

Note that our formulation of a strategy naturally uses a Mealy machine. The output mapping $\lambda_i(.,.)$ has two arguments, the state of the machine and the input s , which lists the outcomes of previous moves within the stage. The transitions take place at the end of the stage. The benefit of using this formulation is that the continuation game is the same at the beginning of each stage. In Chatterjee and Sabourian (2000b), we investigate the effects of modifying this formulation, including studying the effects of having a submachine to play each role. The different formulations can all implement the same strategies, but the complexities in terms of various measures could differ. We refer the reader to that paper for details, but emphasize that in the general unanimity game, the results from other formulations are similar to the one developed here, though they could differ for special cases, like three-player games.

We now consider a machine game, where players first choose machines and then the

machines play the unanimity game in analogy with Abreu-Rubinstein. Using the same lexicographic utility, with complexity coming after bargaining payoffs, what do we find for Nash equilibria of the machine game?

As it turns out, the addition of complexity costs in this setting has some bite but not much. In particular, any division of the pie can be sustained in some Nash equilibrium of the machine game. Perpetual disagreement can, in fact, be sustained by a stationary machine, that is, one that makes the same offers and responses each time, irrespective of past history. Nor can we prove, for general n -player games, that the equilibrium machines will be one state. (A three-player counter-example exists in (Chatterjee and Sabourian 2000b); it does not appear to be possible to generate in games that lasted less than thirty periods.) For two-player games, the result that machines must be one state in equilibrium can be shown neatly (Chatterjee and Sabourian 2000b); another illustration that in this particular area, there is a substantial increase of analytical difficulty in going from two to three players.

One reason why complexity does not appear important here is that the definition of complexity used is too restrictive. Counting the number of states is fine, so long as we don't consider how complex a response might be for partial histories *within* a stage. The next attempt at a solution is based on this observation.

We devise the following definition of complexity: Given the machine and the states, if a machine made the same response to different partial stage histories in different states and another machine made different responses, then the second one was more complex (given that the machines were identical in all other respects). We refer to this notion as *response complexity*. (In (Chatterjee and Sabourian 2000a) the concept of response complexity is in fact stated in terms of the underlying strategy rather than in terms of machines.) It captures the intuition that counting states is not enough; two machines could have the same number of states, for example, because each generated the same number of distinct offers, but the complexity of responses in one machine could be much lower than that in the other. Note that this

notion would only arise in extensive-form games. In normal-form games, counting states could be an adequate measure of complexity. Nor is this notion of complexity derivable from notions of transition complexity, due to Banks and Sundaram, for example, which also apply in normal-form games.

The main result of Chatterjee and Sabourian (2000a) is that this new aspect of complexity enables us to limit the amount of delay that can occur in equilibrium and hence to infer that only one-state machines are equilibrium machines.

The formal proofs using two different approaches are available in Chatterjee and Sabourian (2000a, b). We mention the basic intuition behind these results. Suppose, in the three-player game, there is an agreement in period 4 (this is in the second stage). Why doesn't this agreement take place in period 1 instead? It must be because if the same offer and responses are seen in period 1, some player will reject the offer. But of course, he or she does not have to do so because the required offer never happens. But a strategy that accepts the offer in period 4 and rejects it *off the equilibrium path in period 1 must be more complex, by our definition*, than one that always accepts it whenever it might happen, on or off the expected path. Repeated application of this argument by backward induction gives the result. (The details are more complicated but are in the papers cited above.) Note that this uses the definition that two machines might have the same number of states, and yet one could be simpler than the other. It is interesting, as mentioned earlier, that for two players one can obtain an analogous result without invoking the response simplicity criterion, but from three players on this criterion is essential.

The above result (equilibrium machines have one state each and there are no delays beyond the first stage) is still not enough to refine the set of equilibria to a single allocation. In order to do this, we consider machines that can make errors/trembles in output. As the error goes to zero, we are left with perfect equilibria of our game. With one-state machines, the only subgame perfect equilibria are the ones that give equal division of the pie as $\delta \rightarrow 1$. Thus, a combination of two techniques,

one essentially recognizing that players can make mistakes and the other that players prefer simpler strategies if the payoffs are the same as those given by a more complex strategy, resolves the problem of multiplicity of equilibria in the multi-person bargaining game.

As we mentioned before, the introduction of errors ensures that the equilibrium strategies are credible at every history. We could also take the more direct (and easier) way of obtaining the uniqueness result with complexity costs by considering NEC strategies that are subgame perfect in the underlying game (PEC) (as done in (Chatterjee and Sabourian 2000a)). Then since a history-independent subgame perfect equilibrium of the game is unique and any NEC automaton profile has one state and hence is history independent, it follows immediately that any PEC is unique and induces equal division as $\delta \rightarrow 1$.

Complexity and Repeated Negotiations

In addition to standard repeated games or standard bargaining games, multiplicity of equilibria often appears in dynamic repeated interactions, where a repeated game is superimposed on an alternating offers bargaining game. For instance, consider two firms, in an ongoing vertical relationship, negotiating the terms of a merger. Such situations have been analyzed in several "negotiation models" by Busch and Wen (1995), Fernandez and Glazer (1991), and Haller and Holden (1990). These models can be interpreted as combining the features of both repeated and alternating-offers bargaining games. In each period, one of the two players first makes an offer on how to divide the total available periodic (flow) surplus; if the offer is accepted, the game ends with the players obtaining the corresponding payoffs in the current and every period thereafter. If the offer is rejected, they play some normal-form game to determine their flow payoffs for that period, and then the game moves on to the next period in which the same play continues with the players' bargaining roles reversed. One can think of the normal-form game played in the event of a rejection as a "threat game" in which a player takes actions that could punish the other player by reducing his total payoffs.

If the bargaining had not existed, the game would be a standard repeated normal-form game. Introducing bargaining and the prospect of permanent exit, the negotiation model still admits a large number of equilibria, like standard repeated games. Some of these equilibria involve delay in agreement (even perpetual disagreement) and inefficiency, while some are efficient.

Lee and Sabourian (2007) apply complexity considerations to this model. As in Abreu and Rubinstein (1988) and others, the players choose among automata, and the equilibrium notion is that of NEC and PEC. One important difference however is that in this entry, the authors do not assume the automata to be *finite*. Also, the entry introduces a new machine specification that formally distinguishes between the two *roles* – proposer and responder – played by each player in a given period.

Complexity considerations select only *efficient* equilibria in the negotiation model players are sufficiently patient. First, it is shown that if an agreement occurs in some finite period as an NEC outcome, then it must occur within the first two periods of the game. This is because if an NEC induces an agreement beyond the first two periods, then one of the players must be able to drop the last period's state of his machine without affecting the outcome of the game. Second, given sufficiently patient players, every PEC in the negotiation model that induces perpetual disagreement is at least *long-run* almost efficient; that is, the game must reach a finite date at which the continuation game then on is almost efficient.

Thus, these results take the study of complexity in repeated games a step further from the previous literature in which complexity or bargaining alone has produced only limited selection results. While, as we discussed above, many inefficient equilibria survive complexity refinement, Lee and Sabourian (2007) demonstrate that complexity and bargaining in tandem ensure efficiency in repeated interactions. Complexity considerations also allow Lee and Sabourian to highlight the role of transaction costs in the negotiation game. Transaction costs take the form of paying a cost to enter the bargaining stage of the negotiation game. In contrast to the efficiency result in the

negotiation game with complexity costs, Lee and Sabourian also show that introducing transaction costs into the negotiation game dramatically alters the selection result from efficiency to inefficiency. In particular, they show that, for any discount factor and any transaction cost, every PEC in the costly negotiation game induces perpetual disagreement if the stage game normal form (after any disagreement) has a unique Nash equilibrium.

Complexity, Market Games, and the Competitive Equilibrium

There has been a long tradition in economics of trying to provide a theory of how a competitive market with many buyers and sellers operates. The concept of competitive (Walrasian) equilibrium (see Debreu 1959) is a simple description of such markets. In such an equilibrium each trader chooses rationally the amount he wants to trade taking the prices as given, and the prices are set (or adjust) to ensure that total demanded is equal to the total supplied. The important feature of the setup is that agents assume that they cannot influence (set) the prices, and this is often justified by appealing to the idea that each individual agent is small relative to the market.

There are conceptual as well as technical problems associated with such a justification. First, if no agent can influence the prices, then who sets them? Second, even in a large but finite market, a change in the behavior of a single individual agent may affect the decisions of some others, which in turn might influence the behavior of some other agents and so on and so forth; thus, the market as a whole may end up being affected by the decision of a single individual.

Game theoretic analysis of markets have tried to address these issues (e.g., see Gale 2000; Sabourian 2003). This has turned out to be a difficult task because the strategic analysis of markets, in contrast to the simple and elegant model of competitive equilibrium, tends to be complex and intractable. In particular, dynamic market games have many equilibria, in which a variety of different kinds of behavior are sustained by threats and counter-threats.

More than 60 years ago, Hayek (1945) noted that the competitive markets are simple mechanisms in which economic agents only need to know their own endowments, preferences and technologies, and the vector of prices at which trade takes place. In such environments, economic agents maximizing utility subject to constraints make efficient choices in equilibrium. Below we report some recent work, which suggests that the converse might also be true:

If rational agents have, at least at the margin, an aversion to complex behavior, then their maximizing behavior will result in simple behavioral rules and thereby in a perfectly competitive equilibrium. (Gale and Sabourian 2005)

Homogeneous Markets

In a seminal paper, Rubinstein and Wolinsky (1990), henceforth RW, considered a market for a single indivisible good in which a finite number of homogeneous buyers and homogeneous sellers are matched in pairs and bargain over the terms of trade. In their setup, each seller has one unit of an indivisible good, and each buyer wants to buy at most one unit of the good. Each seller's valuation of the good is 0, and each buyer's valuation is 1. Time is divided into discrete periods, and at each date, buyers and sellers are matched randomly in pairs, and one member of the pair is randomly chosen to be the proposer and the other the responder. In any such match, the proposer offers a price $p \in [0, 1]$ and the responder accepts or rejects the offer. If the offer is accepted, the two agents trade at the agreed price p , and the game ends with the seller receiving a payoff p and the buyer in the trade obtaining a payoff $1 - p$. If the offer is rejected, the pair returns to the market and the process continues. RW further assume that there is no discounting to capture the idea that there is no friction (cost to waiting) in the market.

Assuming that the number of buyers and sellers is not the same, RW showed that this dynamic matching and bargaining game has, in addition to a perfectly competitive outcome, a large set of other subgame perfect equilibrium outcomes, a result reminiscent of the Folk theorem for repeated games. To see the intuition for this, consider the case in which there is one seller

s and many buyers. Since there are more buyers than sellers, the price of 1, at which the seller receives all the surplus, is the unique competitive equilibrium; furthermore, since there are no frictions, $p = 1$ seems to be the most plausible price. RW's precise result, however, establishes that for any price $p^* \in [0, 1]$ and any buyer b^* , there is a subgame perfect equilibrium that results in s and b^* trading at p^* . The idea behind the result is to construct an equilibrium strategy profile such that buyer b^* is identified as the intended recipient of the good at a price p^* . This means that the strategies are such that (i) when s meets b^* , whichever is chosen as the proposer offers price p^* and the responder accepts; (ii) when s is the proposer in a match with some buyer $b \neq b^*$, s offers the good at a price of $p = 1$ and b rejects; and (iii) when a buyer $b \neq b^*$ is the proposer, he offers to buy the good at a price of $p = 0$ and s rejects. These strategies produce the required outcome. Furthermore, the equilibrium strategies make use of the following punishment strategies to deter deviations. If the seller s deviates by proposing to a buyer b a price $p \neq p^*$, b rejects this offer, and the play continues with b becoming the intended recipient of the item at a price of zero. Thus, after rejection by b strategies are the same as those given earlier with the price zero in place of p^* and buyer b in place of buyer b^* . Similarly, if a buyer b deviates by offering a price $p \neq 1$ then the seller rejects, another buyer $b' \neq b$ is chosen to be the intended recipient, and the price at which the unit is traded changes to 1. Further deviations from these punishment strategies can be treated in an exactly similar way.

The strong impression left by RW is that indeterminacy of equilibrium is a robust feature of dynamic market games, and, in particular, there is no reason to expect the outcome to be perfectly competitive. However, the strategies required to support the family of equilibria in RW are quite complex. In particular, when a proposer deviates, the strategies are tailor-made so that the responder is rewarded for rejecting the deviating proposal. This requires coordinating on a large amount of information so that at every information set the players know (and agree) what constitutes a deviation.

In fact, RW show that if the amount of information available to the agents is strictly limited so that the agents do not recall the history of past play, then the only equilibrium outcome is the competitive one. This suggests that the competitive outcome may result if agents use simple strategies. Furthermore, the equilibrium strategies used described in RW to support noncompetitive outcomes are particularly unattractive because they require all players, including those buyers who do not end up trading, to follow complex non-stationary strategies in order to support a noncompetitive outcome. But buyers who do not trade and receive zero payoff on the equilibrium path could always obtain at least zero by following a less complex strategy than the ones specified in RW's construction. Thus, RW's construction of noncompetitive equilibria is not robust if players prefer, at least at the margin, a simpler strategy to a more complex one.

Following the above observation, Sabourian (2003), henceforth S, addresses the role of complexity (simplicity) in sustaining a multiplicity of noncompetitive equilibria in RW's model. The concept of complexity in S is similar to that in Chatterjee and Sabourian (2000a). It is defined by a partial ordering on the set of individual strategies (or automata) that very informally satisfies the following: If two strategies are otherwise identical except that in some role, the second strategy uses more information than that available in the current period of bargaining and the first uses only the information available in the current period, then the second strategy is said to be more complex than the first. S also introduces complexity costs lexicographically into the RW game and shows that any PEC is history-independent and induces the competitive outcome in the sense that all trades take place at the unique competitive price of 1.

Informally, S's conclusions *in the case of a single seller s and many buyers* follow from the following three steps. First, since trading at the competitive price of 1 is the worst outcome for a buyer and the best outcome for the seller, by appealing to complexity type reasoning, it can be shown that in any NEC, a trader's response to a price offer of 1 is always history independent, and thus, he either always rejects 1 or always accepts

1. For example, if in the case of a buyer this was not the case, then since accepting 1 is a worst possible outcome, he could economize on complexity and obtain at least the same payoff by adopting another strategy that is otherwise the same as the equilibrium strategy except that it always rejects 1.

Second, in any noncompetitive NEC in which s receives a payoff of less than 1, there cannot be an agreement at a price of 1 between s and a buyer at *any* history. For example, if at some history a buyer is offered $p = 1$ and he accepts, then by the first step, the buyer should accept $p = 1$ whenever it is offered; but this is a contradiction because it means that the seller can guarantee himself an equilibrium payoff of one by waiting until he has a chance to make a proposal to this buyer.

Third, in any noncompetitive PEC, the continuation payoffs of all buyers are positive at every history. This follows immediately from the previous step because if there is no trade at $p = 1$ at any history, it follows that each buyer can always obtain a positive payoff by offering the seller more than he can obtain in any subgame.

Finally, because of competition between the buyers (there is one seller and many buyers), in any subgame perfect equilibrium, there must be a buyer with a zero continuation payoff after some history. To illustrate the basic intuition for this claim, let m be the worst continuation payoff for s at any history, and suppose that there exists a subgame at which s is the proposer in a match with a buyer $\bar{b}$, and the continuation payoff of s at this subgame is m . Then if at this subgame s proposes $m + \varepsilon$ ($\varepsilon > 0$), $\bar{b}$ must reject (otherwise s can get more than m). Since the total surplus is 1, $\bar{b}$ must obtain at least $1 - m - \varepsilon$ in the continuation game in order to reject s 's offer and s gets at least m ; this implies that the continuation payoff of all $v \neq \bar{v}$ after v 's rejection is less than ε . The result follows by making ε arbitrarily small (and by appealing to the finiteness of f).

But the last two claims contradict each other unless the equilibrium is competitive. This establishes the result for the case in which there is one seller and many buyers. The case of a market with more than one seller is established by induction on the number of sellers.

The matching technology in the above model is random. RW also consider another market game with the matching is endogenous: At each date each seller (the short side of the market) chooses his trading partner. Here, they show that non-competitive outcomes and multiplicity of equilibria survive even when the players discount the future. By strengthening the notion of complexity, S also shows that in the endogenous matching model of RW, the competitive outcome is the only equilibrium if complexity considerations are present.

These results suggest perfectly competitive behavior may result if agents have, at least at the margin, preferences for simple strategies. Unfortunately, both RW and S have too simple a market setup; for example, it is assumed that the buyers are all identical, similarly for the sellers and each agent trades at most one unit of the good. Do the conclusions extend to richer models of trade?

Heterogeneous Markets

There are good reasons to think that it may be too difficult (or even impossible) to establish a similar set of conclusions as in S in a richer framework. For example, consider a heterogeneous market for a single indivisible good, where buyers (and sellers) have a range of valuations of the good and each buyer wants at most one unit of the good and each seller has one unit of the good for sale. In this case the analysis of S will not suffice. First, in the homogeneous market of RW, except for the special case where the number of buyers is equal to the number of sellers, the competitive equilibrium price is either 0 or 1, and all of the surplus goes to one side of the market. S's selection result crucially uses this property of the competitive equilibrium. By contrast, in a heterogeneous market, in general there will be agents receiving positive payoffs on both sides of the market in a competitive equilibrium. Therefore, one cannot justify the competitive outcome simply by focusing on extreme outcomes in which there is no surplus for one party from trade. Second, in a homogeneous market individually rational trade is by definition efficient. This may not be the case in a heterogeneous market (an inefficient trade between inframarginal and an extramarginal

agent can be individually rational). Third, in a homogeneous market, the set of competitive prices remains constant, independently of the set of agents remaining in the market. In the heterogeneous market, this need not be so, and in some cases, the new competitive interval may not even intersect the old one. The change in the competitive interval of prices as the result of trade exacerbates the problems associated with using an induction hypothesis because here future prices may be conditioned on past trades even if prices are restricted to be competitive ones.

Despite these difficulties associated with a market with a heterogeneous set of buyers and sellers, Gale and Sabourian (2005), henceforth GS, show that the conclusions of S can be extended to the case of a heterogeneous market in which each agent trades at most one unit of the good. GS, however, focus on *deterministic* sequential matching models in which one pair of agents is matched at each date and they leave the market if they reach an agreement. In particular, they start by considering exogenous matching processes in which the identities of the proposer and responder at each date are an exogenous and deterministic function of the set of agents remaining in the market and the date. The main result of the entry is that a PEC is always competitive in such a heterogeneous market, thus supporting the view that competitive equilibrium may arise in a finite market where complex behavior is costly.

The notion of complexity in GS is similar to that in Chatterjee and Sabourian (2000a). However, in the GS setup with heterogeneous buyers and sellers, the set of remaining agents changes depending who has traded and left the market and who is remaining, and this affects the market conditions. (In the homogeneous case, only the *number* of remaining agents matters.) Therefore, the definition of complexity in GS is with reference to a given set of remaining agents. GS also discuss an alternative notion of complexity that is independent of the set of remaining agents; such a definition may be too strong and may result in an equilibrium set being empty.

To show their result, GS first establish two very useful restrictions on the strategies that form an

NEC (similar to the no delay result in Chatterjee and Sabourian (2000a)). First, they show that if along the equilibrium path a pair of agents k and ℓ trades at a price p with k as the proposer and ℓ as the responder, then k and ℓ always trade at p , irrespective of the previous history, whenever the two agents are matched in the same way with the same remaining set of agents. To show this consider first the case of the responder ℓ . Then it must be that at every history with the same remaining set of agents, ℓ always accepts p by k . Otherwise, ℓ could economize on complexity by choosing another strategy that is otherwise identical to his equilibrium strategy except that it always accepts p from k without sacrificing any payoff: Such a change of behavior is clearly more simple than sometimes accepting and sometimes rejecting the offer, and moreover, it results in either agent k proposing p and ℓ accepting, so the payoff to agent ℓ is the same as from the equilibrium strategy, or agent k not offering p , in which case the change in the strategy is not observed and the play of the game is unaffected by the deviation. Furthermore, it must also be that at every history with the same remaining set of agents agent k proposes p in any match with ℓ . Otherwise, k could economize on complexity by choosing another strategy that is otherwise identical to his equilibrium strategy except that it always proposes p to ℓ without sacrificing any payoff on the equilibrium path: Such a change of behavior is clearly more simple, and moreover, k 's payoff is not affected because either agents k and ℓ are matched and k proposes p and ℓ by the previous argument accepts, so the payoff to agent k is the same as from the equilibrium strategy, or agents k and ℓ are not matched with k as the proposer, in which case the change in the strategy is not observed and the play of the game is unaffected by the deviation.

GS show a second restriction, again with the same remaining set of agents, namely, that in any NEC for any pair of agents k and ℓ , player ℓ 's response to k 's (on or off-the-equilibrium path) offer is always the same. Otherwise, it follows that ℓ sometimes accepts an offer p by k and sometimes rejects (with the same remaining set of agents). Then by the first restriction it must be

that if such an offer is made by k to ℓ on the equilibrium path it is rejected. But then ℓ can could economize on complexity by always rejecting p by k without sacrificing any payoff on the equilibrium path: Such a change of behavior is clearly more simple, and furthermore, ℓ 's payoff is not affected because such a behavior is the same as what the equilibrium strategy prescribes on the equilibrium path.

By appealing to the above two properties of NEC and to the competitive nature of the market GS establish, using a complicated induction argument, that every PEC induces a competitive outcome in which each trade occurs at the same competitive price.

The matching model we have described so far is deterministic and exogenous. The selection result of GS however extends to richer deterministic matching models. In particular, GS also consider a semi-endogenous sequential matching model in which the choice of partners is endogenous but the identity of the proposer at any date is exogenous. Their results extend to this variation, with an endogenous choice of responders. A more radical departure change would be to consider the case where at any date any agent can choose his partner and make a proposal. Such a totally endogenous model of trade generates new conceptual problems. In a recent working paper, Gale and Sabourian (2008) consider a continuous time version of such a matching model and show that complexity considerations allow one to select a competitive outcome in the case of totally endogenous matching. Since the selection result holds for all the different matching models, we can conclude that complexity considerations inducing a competitive outcome seem to be a robust result in deterministic matching and bargaining market games with heterogeneous agents.

Random matching is commonly used in economic models because of its tractability. The basic framework of GS, however, does not extend to such a framework if either the buyers or the sellers are not identical. This is for two different reasons. First, in general in any random framework, there is more than one outcome path that can occur in equilibrium with a positive probability; as a result introducing complexity lexicographically may not

be enough to induce agents to behave in a simple way (they will have to be complex enough to play optimally along all paths that occur with a positive probability). Second, in Gale and Sabourian (2006) it is shown that subgame perfect equilibria in Markov strategies are not necessarily perfectly competitive for the random matching model with heterogeneous agents. Since the definition of complexity in GS is such that Markov strategies are the least complex ones, it follows that with random matching the complexity definition used in GS is not sufficient to select a competitive outcome.

Complexity and Off-The-Equilibrium-Path Play

The concept of the PEC (or NEC) used in S, GS, and elsewhere was defined to be such that for each player, the strategy/automaton has minimal complexity among all strategies/automata that are best responses to the equilibrium strategies/automata of others. Although, these concepts are very mild in the treatment of complexity, it should be noted that there are other ways of introducing complexity into the equilibrium concept. One extension of the above setup is to treat complexity as a (small) positive fixed cost of choosing a more complex strategy and define a Nash (subgame perfect) equilibrium with a fixed positive complexity costs accordingly. All the selection results based on lexicographic complexity in the papers we discuss in this survey also hold for positive small complexity costs. This is not surprising because with positive costs complexity has at least as much bite as in the lexicographic case; there is at least as much refinement of the equilibrium concept with the former as with the latter. In particular, in the case of an NEC (or a PEC), *in considering complexity*, players ignore any consideration of payoffs off the equilibrium path, and the trade-off is between the equilibrium payoffs of two strategies and the complexity of the two. As a result these concepts put more weight on complexity costs than on being “prepared” for off-the-equilibrium-path moves. Therefore, although complexity costs are insignificant, they take priority over optimal behavior after deviations. (See Chatterjee and Sabourian 2000b for a discussion.)

A different approach would be to assume that complexity is a less significant criterion than the off-the-equilibrium payoffs. In the extreme case, one would require agents to choose minimally complex strategies among the set of strategies that are best responses on and off the equilibrium path (see Kalai and Neme 1992).

An alternative way of illustrating the differences between the different approaches is by introducing two kinds of vanishingly small perturbations into the underlying game. One perturbation is to impose a small but positive cost of choosing a more complex strategy. Another perturbation is to introduce a small but positive probability of making an error (off-the-equilibrium-path move). Since a PEC requires each agents to choose a minimally complex strategy within the set of best responses, it follows that the limit points of Nash equilibria of the above perturbed game correspond to the concept of PEC if we first let the probability of making an off-the-equilibrium-path move go to zero and then let the cost of choosing a more complex strategy go to zero (this is what Chatterjee and Sabourian (2000a) do). On the other hand, in terms of the above limiting arguments, if we let the cost of choosing a more complex strategy go to zero and then let the probability of making an off-the-equilibrium-path move go to zero, then any limit corresponds to the equilibrium definition in Kalai and Neme (1992) where agents choose minimally complex strategies among the set of strategies that are best responses on and off the equilibrium path.

Most of the results reported in this entry on refinement and endogenous complexity (e.g., Abreu-Rubinstein (1988)), Chatterjee and Sabourian (2000a), Gale and Sabourian (2005), and Lee and Sabourian (2007) hold only for the concept of NEC and its variations and thus depend crucially on assuming that complexity costs are more important than off-the-equilibrium payoffs. This is because these results always appeal to an argument that involves economizing on complexity if the complexity is not used off the equilibrium path. Therefore, they may be a good predictor of what may happen only if complexity costs are more significant than the perturbations that induce off-the-equilibrium-path behavior.

The one exception is the selection result in S (Sabourian 2003). Here, although the result we have reported is stated for NEC and its variations, it turns out that the selection of competitive equilibrium does not in fact depend on the relative importance of complexity costs and off-the-equilibrium-path payoffs. It remains true even for the case where the strategies are required to be least complex among those that are best responses at every information set. This is because in S's analysis complexity is only used to show that every agent's response to the price offer of 1 is always the same irrespective of the past history of play. This conclusion holds irrespective of the relative importance of complexity costs and off-the-equilibrium payoff because trading at the price of 1 is the best outcome that any seller can achieve at any information set (including those off-the-equilibrium) and a worst outcome for any buyer. Therefore, irrespective of the order, the strategy of sometimes accepting a price of 1 and sometimes rejecting cannot be an equilibrium for a buyer (similar arguments applies for a seller) because the buyer can economize on complexity by always rejecting the offer without sacrificing any payoff off- or on-the-equilibrium path (accepting $p = 1$ is a worse possible outcome).

Discussion and Future Directions

The use of finite automata as a model of players in a game has been criticized as being inadequate, especially because as the number of states becomes lower it becomes more and more difficult for the small automaton to do routine calculations, let alone the best response calculations necessary for game-theoretic equilibria. Some of the papers we have explored address other aspects of complexity that arise from the concrete nature of the games under consideration. Alternative models of complexity are also suggested, such as computational complexity and communication complexity.

While our work and the earlier work on which it builds focus on equilibrium, an alternative approach might seek to see whether simplicity evolves in some reasonable learning model.

Maenner (2008) has undertaken such an investigation with the infinitely repeated Prisoners' Dilemma (studied in the equilibrium context by Abreu and Rubinstein). Maenner provides an argument for "learning to be simple." On the other hand, there are arguments for increasing complexity in competitive games (Robson 2003). It is an open question, therefore, whether simplicity could arise endogenously through learning, though it seems to be a feature of most human preferences and aesthetics (see Birkhoff 1933).

The broader research program of explicitly considering complexity in economic settings might be a very fruitful one. Auction mechanisms are designed with an eye towards how complex they are – simplicity is a desideratum. The complexity of contracting has given rise to a whole literature on incomplete contracts, where some models postulate a fixed cost per contingency described in the contract. All this is apart from the popular literature on complexity, which seeks to understand complex, adaptive systems from biology. The use of formal complexity measures such as those considered in this survey and the research we describe might throw some light on whether incompleteness of contracts, or simplicity of mechanisms, is an assumption or a result (of explicitly considering choice of level of complexity).

Acknowledgments We wish to thank an anonymous referee and Jihong Lee for valuable comments that improved the exposition of this entry. We would also like to thank St. John's College, Cambridge, and the Pennsylvania State University for funding Dr. Chatterjee's stay in Cambridge at the time this entry was written.

Bibliography

- Abreu D, Rubinstein A (1988) The structure of Nash equilibria in repeated games with finite automata. *Econometrica* 56:1259–1282
- Anderlini L (1990) Some notes on church's thesis and the theory of games. *Theory Decis* 29:19–52
- Anderlini L, Sabourian H (1995) Cooperation and effective computability. *Econometrica* 63:1337–1369
- Aumann RJ (1981) Survey of repeated games. In: *Essays in game theory and mathematical economics in honor of*

- Oskar Morgenstern. Bibliographisches Institut, Mannheim/Vienna/Zurich, pp 11–42
- Banks J, Sundaram R (1990) Repeated games, finite automata and complexity. *Games Econ Behav* 2:97–117
- Ben Porath E (1986) Repeated games with bounded complexity. Mimeo, Stanford University, Stanford, Calif
- Ben Porath E (1993) Repeated games with finite automata. *J Econ Theory* 59:17–32
- Binmore KG (1987) Modelling rational players I. *Econ Philos* 3:179–214
- Binmore KG, Samuelson L (1992) Evolutionary stability in repeated games played by finite automata. *J Econ Theory* 57:278–305
- Binmore KG, Piccione M, Samuelson L (1998) Evolutionary stability in alternating-offers bargaining games. *J Econ Theory* 80:257–291
- Birkhoff GD (1933) Aesthetic measure. Harvard University Press, Cambridge, MA
- Bloise G (1998) Strategic complexity and equilibrium in repeated games. Unpublished doctoral dissertation, University of Cambridge
- Busch L-A, Wen Q (1995) Perfect equilibria in a negotiation model. *Econometrica* 63:545–565
- Chatterjee K (2002) Complexity of strategies and multiplicity of Nash equilibria. *Group Decis Negot* 11:223–230
- Chatterjee K, Sabourian H (2000a) Multiperson bargaining and strategic complexity. *Econometrica* 68:1491–1509
- Chatterjee K, Sabourian H (2000b) N-person bargaining and strategic complexity. Mimeo, University of Cambridge and the Pennsylvania State University, Cambridge, UK and University Park, Pa., USA
- Debreu G (1959) Theory of value. Yale University Press, New Haven/London
- Fernandez R, Glazer J (1991) Striking for a bargain between two completely informed agents. *Am Econ Rev* 81:240–252
- Fudenberg D, Maskin E (1990) Evolution and repeated games. Mimeo, Harvard University, Cambridge, Mass
- Fudenberg D, Tirole J (1991) Game theory. MIT Press, Cambridge, MA
- Gale D (2000) Strategic foundations of general equilibrium: dynamic matching and bargaining games. Cambridge University Press, Cambridge
- Gale D, Sabourian H (2005) Complexity and competition. *Econometrica* 73:739–770
- Gale D, Sabourian H (2006) Markov equilibria in dynamic matching and bargaining games. *Games Econ Behav* 54:336–352
- Gale D, Sabourian H (2008) Complexity and competition II; endogenous matching. Mimeo, New York University, New York, USA/University of Cambridge, Cambridge, UK
- Haller H, Holden S (1990) A letter to the editor on wage bargaining. *J Econ Theory* 52:232–236
- Hayek F (1945) The use of knowledge in society. *Am Econ Rev* 35:519–530
- Herrero M (1985) A strategic theory of market institutions. Unpublished doctoral dissertation, London School of Economics
- Kalai E, Neme A (1992) The strength of a little perfection. *Int J Game Theory* 20:335–355
- Kalai E, Stanford W (1988) Finite rationality and interpersonal complexity in repeated games. *Econometrica* 56:397–410
- Klemperer P (ed) (2000) The economic theory of auctions. Elgar, Northampton
- Lee J, Sabourian H (2007) Coase theorem, complexity and transaction costs. *J Econ Theory* 135:214–235
- Maenner E (2008) Adaptation and complexity in repeated games. *Games Econ Behav* 63:166–187
- Miller GA (1956) The magical number seven plus or minus two: Some limits on our capacity to process information. *Psychol Rev* 63:81–97
- Neme A, Quintas L (1995) Subgame perfect equilibrium of repeated games with implementation cost. *J Econ Theory* 66:599–608
- Neyman A (1985) Bounded complexity justifies cooperation in the finitely-repeated Prisoners' Dilemma. *Econ Lett* 19:227–229
- Neyman A (1997) Cooperation, repetition and automata in cooperation: game-theoretic approaches. In: Hart S, Mas-Colell A (eds) NATO ASI series F, vol 155. Springer, Berlin, pp 233–255
- Osborne M, Rubinstein A (1990) Bargaining and markets. Academic, New York
- Osborne M, Rubinstein A (1994) A course in game theory. MIT Press, Cambridge, MA
- Papadimitriou CH (1992) On games with a bounded number of states. *Games Econ Behav* 4:122–131
- Piccione M (1992) Finite automata equilibria with discounting. *J Econ Theory* 56:180–193
- Piccione M, Rubinstein A (1993) Finite automata play a repeated extensive game. *J Econ Theory* 61:160–168
- Robson A (2003) The evolution of rationality and the Red Queen. *J Econ Theory* 111:1–22
- Rubinstein A (1982) Perfect equilibrium in a bargaining model. *Econometrica* 50:97–109
- Rubinstein A (1986) Finite automata play the repeated Prisoners' Dilemma. *J Econ Theory* 39:83–96
- Rubinstein A (1998) Modeling bounded rationality. MIT Press, Cambridge, MA
- Rubinstein A, Wolinsky A (1990) Decentralized trading, strategic behaviour and the Walrasian outcome. *Rev Econ Stud* 57:63–78
- Sabourian H (2003) Bargaining and markets: complexity and the competitive outcome. *J Econ Theory* 116:189–228
- Selten R (1965) Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit. *Z gesamte Staatswiss* 12:201–324
- Shaked A (1986) A three-person unanimity game. In: The Los Angeles national meetings of the Institute of Management Sciences and the Operations Research Society of America, Mimeo, University of Bonn, Bonn, Germany
- Zemel E (1989) Small talk and cooperation: a note on bounded rationality. *J Econ Theory* 49:1–9

Part II

Agent-Based Models

Agent-Based Modeling and Simulation, Introduction to

Filippo Castiglione

Istituto Applicazioni del Calcolo (IAC), Consiglio Nazionale delle Ricerche (CNR), Rome, Italy

Keywords

Discrete mathematical modeling · Simulation · Complex systems

Agent-based modeling (ABM) is a computational modeling paradigm that is markedly useful in studying complex systems composed of a large number of interacting entities with many degrees of freedom. Other names for ABM are individual-based modeling (IBM) or multi-agent systems (MAS). Physicists often use the term micro-simulation or interaction-based computing.

The basic idea of ABM is to construct the computational counterpart of a conceptual model of a system under study on the basis of discrete entities (agents) with defined properties and behavioral rules and then to simulate them in a computer to mimic the real phenomena.

The definition of agent is somewhat fuzzy as witnessed by the fact that the models found in the literature adopt an extremely heterogeneous rationale. The agent is an autonomous entity having its own internal state reflecting its perception of the environment and interacting with other entities according to more or less sophisticated rules. In practice, the term agent is used to indicate entities ranging all the way from simple pieces of software to “conscious” entities with learning capabilities. For example, there are “helper” agents for web retrieval, robotic agents to explore inhospitable environments, buyer/seller agents in an economy, and so on. Roughly speaking, an entity is an “agent” if it has some degree of autonomy, that is, if it is distinguishable from its environment by some kind of spatial, temporal, or functional attribute: an agent must be identifiable. Moreover, it is usually required that an agent must have some

autonomy of action and that it must be able to engage in tasks in an environment without direct external control.

From simple agents, which interact locally with simple rules of behavior, merely responding befittingly to environmental cues, and not necessarily striving for an overall goal, we observe a synergy which leads to a higher-level whole with much more intricate behavior than the component agents (*holism*, meaning all, entire, total). Agents can be identified on the basis of a set of properties that must characterize an entity and, in particular, *autonomy* (the capability of operating without intervention by humans, and a certain degree of control over its own state); *social ability* (the capability of interacting by employing some kind of agent communication language); *reactivity* (the ability to perceive an environment in which it is situated and respond to perceived changes); and *pro-activeness* (the ability to take the initiative, starting some activity according to internal goals rather than as a reaction to an external stimulus). Moreover, it is also conceptually important to define what the agent “environment” in an ABM is.

In general, given the relative immaturity of this modeling paradigm and the broad spectrum of disciplines in which it is applied, a clear-cut and widely accepted definition of high-level concepts of agents, environment, interactions, and so on is still lacking. Therefore a real ABM ontology is needed to address the epistemological issues related to the agent-based paradigm of modeling of complex systems in order to attempt to reach a more general comprehension of emergent properties which, though ascribed to the definition of a specific application domain, are also universal (see chapter ► “Agent-Based Modeling and Simulation”).

Historically, the first simple conceptual form of agent-based models was developed in the late 1940s, and it took the advent of the computer to show its modeling power. This is the Von Neumann machine, a theoretical machine capable of reproduction. The device von Neumann proposed would follow precisely detailed instructions to

produce an identical copy of itself. The concept was then improved by Stanislaw Ulam. He suggested that the machine be built on paper, as a collection of cells on a grid. This idea inspired von Neumann to create the first of the models later termed cellular automata (CA). John Conway then constructed the well-known “Game of Life.” Unlike the von Neumann’s machine, Conway’s Game of Life operated by simple rules in a virtual world in the form of a two-dimensional checkerboard.

Conway’s Game of Life has become a paradigmatic example of models concerned with the emergence of order in nature. How do systems self-organize themselves and spontaneously achieve a higher-ordered state? These and other questions have been deeply addressed in the first workshop on Artificial Life (ALife) held in the late 1980s in Santa Fe. This workshop shaped the ALife field of research. Agent-based modeling is historically connected to ALife because it has become a distinctive form of modeling and simulation in this field. In fact, the essential features of ALife models are translated into computational algorithms through agent-based modeling (see chapter ► [“Agent-Based Modeling and Artificial Life”](#)).

Agent-based models can be seen as the natural extension of the CA-like models, which have been very successful in the past decades in shedding light on various physical phenomena. One important characteristic of ABMs, which distinguishes them from cellular automata, is the potential asynchrony of the interactions among agents and between agents and their environments. In ABMs, agents typically do not simultaneously perform actions at constant time-steps, as in CAs or Boolean networks. Rather, their actions follow discrete-event cues or a sequential schedule of interactions. The discrete-event setup allows for the cohabitation of agents with different environmental experiences. Also ABMs are not necessarily grid based nor do agents “tile” the environment.

Physics investigation is based on building models of reality. It is a common experience that, even using simple “building blocks,” one usually obtains systems whose behavior is quite complex. This is the reason why CA-like, and therefore agent-based models, have been used extensively among physicists to investigate

experimentally (that is, on a computer) the essential ingredients of a complex phenomenon. Rather than being derived from some fundamental law of physics, these essential ingredients constitute artificial worlds. Therefore, there exists a pathway from Newton’s laws to CA and ABM simulations in classical physics that has not yet expressed all its potential (see chapter ► [“Interaction-Based Computing in Physics”](#)).

CA-like models also proved very successful in theoretical biology to describe the aggregation of cells or microorganisms in normal or pathological conditions (see chapter ► [“Cellular Automaton Modeling of Tumor Invasion”](#)).

Returning to the concept of agent in the ABM paradigm, an agent may represent a particle, a financial trader, a cell in a living organism, a predator in a complex ecosystem, a power plant, an atom belonging to a certain material, a buyer in a closed economy, customers in a market model, forest trees, cars in large traffic vehicle system, etc. Once the level of description of the system under study has been defined, the identification of such entities is quite straightforward. For example, if one looks at the world economy, then the correct choice of agents are nations, rather than individual companies. On the other hand, if one is interested in looking at the dynamics of a stock, then the entities determining the price evolution are the buyers and sellers.

This example points to a field where ABM provides a very interesting and valuable instrument of research. Indeed, mainstream economic models typically make the assumption that an entire group of agents, for example, “investors,” can be modeled with a single “rational representative agent.” While this assumption has proven extremely useful in advancing the science of economics by yielding analytically tractable models, it is clear that the assumption is not realistic: people differ in their tastes, beliefs, and sophistication, and as many psychological studies have shown, they often deviate from rationality in systematic ways. Agent-based computational economics (ACE) is a framework allowing economics to expand beyond the realm of the “rational representative agent.” By modeling and simulating the behavior of each agent and interactions among

agents, agent-based simulation allows us to investigate the dynamics of complex economic systems with many heterogeneous (and not necessarily fully rational) agents. Agent-based computational economics complements the traditional analytical approach and is gradually becoming a standard tool in economic analysis (see chapter ► [“Agent-Based Computational Economics”](#)).

Because the paradigm of agent-based modeling and simulation can handle richness of detail in the agent’s description and behavior, this methodology is very appealing for the study and simulation of *social systems*, where the behavior and the heterogeneity of the interacting components are not safely reducible to some stylized or simple mechanism. Social phenomena simulation in the area of agent-based modeling and simulation, concerns the emulation of the individual behavior of a group of social entities, typically including their cognition, actions, and interaction. This field of research aims at “growing” artificial society following a bottom-up approach.

Historically, the birth of the agent-based model as a model for social systems can be primarily attributed to a computer scientist, Craig Reynolds. He tried to model the reality of lively biological agents, known as artificial life, a term coined by Christopher Langton. In 1996 Joshua M. Epstein and Robert Axtell developed the first large-scale agent model, the Sugarscape, to simulate and explore the role of social phenomenon such as seasonal migrations, pollution, sexual reproduction, combat, transmission of disease, and even culture (see chapter ► [“Social Phenomena Simulation”](#)).

In the field of artificial intelligence, the collective behavior of agents that without central control, collectively carry out tasks normally requiring some form of “intelligence,” constitutes the central concept in the field of *swarm intelligence*. The term “swarm intelligence” first appeared in 1989. As the use of the term swarm intelligence has increased, its meaning has broadened to a point in which it is often understood to encompass almost any type of collective behavior. Technologically, the importance of “swarms” is mainly based on potential advantages over centralized systems. The potential advantages are economy (the swarm units are simple, hence, in principle, mass producible, *modularizable*,

interchangeable, and disposable); reliability (due to the redundancy of the components; destruction/death of some units has negligible effect on the accomplishment of the task, as the swarm adapts to the loss of few units); and ability to perform tasks beyond those of centralized systems, for example, escaping enemy detection. From this initial perspective on potential advantages, the actual application of swarm intelligence has extended to many areas and inspired potential future applications in defense and space technologies (for example, control of groups of unmanned vehicles in land, water, or air), flexible manufacturing systems, and advanced computer technologies (bio-computing), medical technologies, and telecommunications (see chapter ► [“Swarm Intelligence”](#)).

Similarly, robotics has adopted the ABM paradigm to study, by means of simulation, the crucial features of adaptation and cooperation in the pursuit of a global goal. Adaptive behavior concerns the study of how organisms develop their behavioral and cognitive skills through a synthetic methodology, consisting of designing artificial agents which are able to adapt to their environment autonomously. These studies are important both from a modeling point of view (that is, for better understanding intelligence and adaptation in natural beings) and from an engineering point of view (that is, for developing artifacts displaying effective behavioral and cognitive skills) (see chapter ► [“Embodied and Situated Agents, Adaptive Behavior in”](#)).

What makes ABM a novel and interesting paradigm of modeling is the idea that agents are individually represented and “monitored” in the computer’s memory. One can, at any time during the simulation, ask a question such as “what is the age distribution of the agents?” or “how many stocks have accumulated buyers following that specific strategy?” or “what is the average velocity of the particles?” “Large-scale” simulations in the context of agent-based modeling are not only simulations that are large in terms of size (number of agents simulated) but are also complex. Complexity is inherent in agent-based models, as they are usually composed of dynamic, heterogeneous, interacting agents. Large-scale agent-based models have also been referred to as

“Massively Multi-agent Systems (MMAS).” MMAS is defined as “beyond resource limitation”: the number of agents exceeds local computer resources, or the situations are too complex to design and program given human cognitive resource limits. Therefore, for agent-based modeling, “large scale” is not simply a size problem, it is also a problem of managing complexity to ensure scalability of the agent model. Agent-based models increase in scale as the modeler requires many agents to investigate whole system behavior, or as the modeler wishes to fully examine the response of a single agent in a realistic context. There are two key problems that have to be tackled as the scale of a multi-agent system increases: computational resources limit the simulation time and/or data storage capacity, and agent model analysis may become more difficult. Difficulty in analyzing the model may be due to the model system having a large number of complex components or due to memory for model output storage being restricted by computer resources (see chapter ► [“Agent-Based Modeling, Large-Scale Simulations”](#)).

For implementation of agent-based models, both domain-specific and general-purpose languages are routinely used. Domain-specific languages include business-oriented languages (for example, spreadsheet programming tools); science and engineering languages (such as Mathematica); and dedicated agent-based modeling languages (for example, NetLogo or Swarm). General-purpose languages can be used directly (as in the case of Java programming) or within agent-based modeling toolkits (for example, Repast). The choice that is most appropriate for a given modeling project depends on both the requirements of that project and the resources available to implement it (see chapter ► [“Agent-Based Modeling and Computer Languages”](#)).

Interestingly, ABM is not being used exclusively in science. In fact, the entertainment industry has promoted its own interest in the ABM technology. As graphics technology has improved in recent years, more and more importance has been placed on the behavior of virtual characters in applications set in virtual worlds in areas such as games, movies, and simulations. The behavior

of virtual characters should be believable in order to create the illusion that these virtual worlds are populated with living characters. This has led to the application of agent-based modeling to the control of these virtual characters. There are a number of advantages of using agent-based modeling techniques which include the fact that they remove the requirement for hand controlling all agents in a virtual environment and allow agents in games to respond to unexpected actions by players (see chapter ► [“Computer Graphics and Games, Agent-Based Modeling in”](#)).

Since it is difficult to formally analyze complex multi-agent systems, they are mainly studied through computer simulations. While computer simulations can be very useful, results obtained through simulations do not formally validate the observed behavior. It is widely recognized that building a sound and widely applicable theory for ABM systems will require an interdisciplinary approach and the development of new mathematical and computational concepts. In other words, there is a compelling need for a mathematical framework which one can use to represent ABM systems and formally establish their properties. Fortunately, some known mathematical frameworks already exist that can be used to formally describe multi-agent systems, for example, that of finite *dynamical systems* (both deterministic and stochastic). A sampling of the results from this field of mathematics shows that they can be used to carry out rigorous studies of the properties of multi-agent systems and, in general, that they can also serve as universal models for computation. Moreover, special cases of dynamical systems (sequential dynamical systems) can be structured in accordance with the theory of categories and therefore provide the basis for a formal theory to describe ABM behavior (see chapter ► [“Agent-Based Modeling, Mathematical Formalism for”](#)).

On the same line of thought, agents and interaction can be studied from the perspective of logic and computer science. In particular, ideas about *logical dynamics*, *game semantics*, and *geometry of interaction*, which have been developed over the past two decades, lead towards a structural theory of agents and interaction. This provides a basis for powerful logical methods such as

compositionality, types and high-order calculi, which have proved so fruitful in computer science, to be applied in the domain of ABM and simulation (see chapter ► [“Logic and Geometry of Agents in Agent-Based Modeling”](#)).

The appeal of ABM methodology in science increases manifestly with advances in computational power of modern computers. However, it is important to bear in mind that increasing the complexity of a model does not necessarily bring more understanding of the fundamental laws governing the overall dynamics. Actually,

beyond a certain level of model complexity, the model loses its ability to explain or predict reality and it reduces to a mere surrogate of reality where things happen with a surprisingly good adherence to reality, although we are unable to explain *why* this happens. Therefore, model construction must proceed incrementally, step by step, possibly validating the model at each stage of development before adding more details. ABM technology is very powerful but, if badly used, could reduce science to a mere exercise consisting of mimicking reality.

Agent-Based Modeling and Simulation

Stefania Bandini, Sara Manzoni and
Giuseppe Vizzari
Complex Systems and Artificial Intelligence
Research Center, University of Milan-Bicocca,
Milan, Italy

Article Outline

Glossary
Definition of the Subject
Introduction
Agent-Based Models for Simulation
Platforms for Agent-Based Simulation
Future Directions
Bibliography

Glossary

Agent The definition of the term agent is controversial even inside the restricted community of computer scientists dealing with research on agent models and technologies (Franklin and Graesser 1997). A weak definition, that could be suited to describe the extremely heterogeneous approaches in the agent-based simulation context, is “an autonomous entity, having the ability to decide the actions to be carried out in the environment and interactions to be established with other agents, according to its perceptions and internal state”.

Agent architecture The term agent *architecture* (Russel and Norvig 1995) refers to the internal structure that is responsible of effectively selecting the actions to be carried out, according to the perceptions and internal state of an agent. Different architectures have been proposed in order to obtain specific agent behaviors and they are generally classified into *deliberative* and *reactive* (respectively, *hysteretic* and

tropicistic, according to the classification reported in Genesereth and Nilsson (1987)).

Autonomy The term autonomy has different meanings, for it represents (in addition to the control of an agent over its own internal state) different aspects of the possibility of an agent to decide about its own actions. For instance, it may represent the possibility of an agent to decide (i) about the timing of an action, (ii) whether or not to fulfill a request, (iii) to act without the need of an external trigger event (also called pro-activeness or pro-activity) or even (iv) basing on its personal experience instead of hard-wired knowledge (Russel and Norvig 1995). It must be noted that different agent models do not generally embody all the above notions of autonomy.

Interaction “An interaction occurs when two or more agents are brought into a dynamic relationship through a set of reciprocal actions” (Ferber 1999).

Environment “The environment is a first-class abstraction that provides the surrounding conditions for agents to exist and that mediates both the interaction among agents and the access to resources” (Weyns et al. 2007).

Platform for agent-based simulation a software framework specifically aimed at supporting the realization of agent-based simulation systems; this kind of framework often provides abstractions and mechanisms for the definition of agents and their environments, to support their interaction, but also additional functionalities like the management of the simulation (e.g., setup, configuration, turn management), its visualization, monitoring and the acquisition of data about the simulated dynamics.

Definition of the Subject

Agent-Based Modeling and Simulation – an approach to the modeling and simulation of a system in which the overall behavior is determined by

the local action and interaction of a set of agents situated in an environment. Every agent chooses the action to be carried out on the basis of its own behavioral specification, internal state and perception of the environment. The environment, besides enabling perceptions, can regulate agents' interactions and constraint their actions.


Introduction

Computer simulation represents a way to exploit a computational model to evaluate designs and plans without actually bringing them into existence in the real world (e.g., architectural designs, road networks and traffic lights), but also to evaluate theories and models of complex systems (e.g., biological or social systems) by envisioning the effect of the modeling choices, with the aim of gaining insight of their functioning. The use of these "synthetic environments" is sometimes necessary because the simulated system cannot actually be observed, since it is actually being designed or also for ethical or practical reasons. A general schema (based on several elaborations, such as those described in Edmonds (2001) and Gilbert and Troitzsch (2005)) describing the role of simulation as a predictive or explanatory instrument is shown in Fig. 1.

Several situations are characterized by the presence of autonomous entities whose actions and interactions determine (in a non-trivial way) the evolution of the overall system. Agent-based models are particularly suited to represent these situations, and to support the study and analysis of

topics like decentralized decision making, local-global interactions, self-organization, emergence, effects of heterogeneity in the simulated system. The interest in this relatively recent approach to modeling and simulation is demonstrated by the number of scientific events focused in this topic (see, to make some examples rooted in the computer science context, the Multi Agent Based Simulation workshop series (Davidsson et al. 2005; Hales et al. 2003; Moss and Davidsson 2001; Sichman and Antunes 2006; Sichman et al. 1998; Sichman et al. 2003), the IMA workshop on agent-based modeling (<http://www.ima.umn.edu/complex/fall/agent.html>) and the Agent-Based Modeling and Simulation symposium (Bandini et al. 2006a)). Agent-based models and multi-agent systems (MAS) have been adopted to simulate complex systems in very different contexts, ranging from social and economical simulation (see, e.g., Dosi et al. 2006) to logistics optimization (see, e.g., Weyns et al. 2006b), from biological systems (see, e.g., Bandini et al. 2006b) to traffic (see, e.g., Balmer and Nagel 2006; Bazzan et al. 1999; Wahle and Schreckenberg 2001) and crowd simulation (see, e.g., Batty 2001).

This heterogeneity in the application domains also reflects the fact that, especially in this context of agent focused research, influences come from most different research areas. Several *traffic and crowd* agent models, to make a relevant example, are deeply influenced by *physics*, and the related models provide agents that are modeled as particles subject to forces generated by the environment as



Agent-Based Modeling and Simulation, Fig. 1 A general schema describing the usage of simulation as a predictive or explanatory instrument

well as by other agents (i.e., *active walker* models, such as Helbing et al. 1997). Other approaches to crowd modeling and simulation build on experiences with *Cellular Automata* (CA) approaches (see, e.g., Schadschneider et al. 2002) but provide a more clear separation between the environment and the entities that inhabit, act and interact in it (see, e.g., Bandini et al. 2004; Henein and White 2005). This line of research leads to the definition of models for situated MASs, a type of model that was also defined and successfully applied in the context of (reactive) robotics and control systems (Weyns and Holvoet 2006; Weyns et al. 2005). Models and simulators defined and developed in the context of social sciences (Gilbert and Troitzsch 2005) and economy (Pyka and Fagiolo 2007) are instead based on different theories (often non-classical ones) of *human behavior* in order to gain further insight on it and help building and validating new theories.

The common standpoint of all the above-mentioned approaches and of many other ones that describe themselves as agent-based is the fact that the analytical unit of the system is represented by the individual agent, acting and interacting with other entities in a shared environment: the overall system dynamic is not defined in terms of a global function, but rather the result of individuals' actions and interactions. On the other hand, it must also be noted that in most of the

introduced application domains, the environment plays a prominent role because:

- It deeply influences the behaviors of the simulated entities, in terms of perceptions and allowed actions for the agents.
- The aim of the simulation is to observe some aggregate level behavior (e.g., the density of a certain type of agent in an area of the environment, the average length of a given path for mobile agents, the generation of clusters of agents), that can actually only be observed in the environment.


Besides these common elements, the above introduced approaches often dramatically differ in the way agents are described, both in terms of properties and behavior. A similar consideration can be done for their environment.

Considering the above introduced considerations, the aim of this article is not to present a specific technical contribution but rather to introduce an abstract reference model that can be applied to analyze, describe and discuss different models, concrete simulation experiences, platforms that, from different points of view, legitimately claim to adopt an agent-based approach. The reference model is illustrated in Fig. 2.

In particular, the following are the main elements of this reference model:

Agent-Based Modeling and Simulation,

Fig. 2 An abstract reference model to analyze, describe and discuss different models, concrete simulation experiences, platforms legitimately claiming to adopt an agent-based approach



- *Agents* encompassing a possibly heterogeneous behavioral specification;
- Their *environment*, supplying agents their perceptions and enabling their actions;
- Mechanisms of *interaction* among agents.

Agent interaction, in fact, can be considered as a specific type of action having a central role in agent-based models. In the following Section, these elements will be analyzed, with reference to the literature on agent-based models and technologies but also considering the concrete experiences and results obtained by researches in other areas, from biology, to urban planning, to social sciences, in an attempt to enhance the mutual awareness of the points of contact of the respective efforts. Section “Platforms for Agent-Based Simulation” briefly discusses the available platforms supporting the rapid prototyping or development of agent-based simulations. A discussion on the future directions for this multi-disciplinary research area will end the article.

Agent-Based Models for Simulation

A model is an abstract and simplified representation of a given part of the world, either existing or planned (a target system, in the terminology adopted for Fig. 1). Models are commonly defined in order to study and explain observed phenomena or to forecast future phenomena. Agent-based models for simulation, as previously mentioned, are characterized by the presence of agents performing some kind of behavior in a shared environment. The notion of agent, however, is controversial even inside the restricted community of computer scientists dealing with research on agent models and technologies (Franklin and Graesser 1997). The most commonly adopted definition of agent (Wooldridge and Jennings 1995) specifies a set of properties that must characterize an entity to effectively call it an agent, and in particular *autonomy* (the possibility to operate without intervention by humans, and a certain degree of control over its own state), *social ability* (the possibility to interact employing some kind of agent communication language, a notion that will

be analyzed more in details in Subsect. “Agent Interaction”), *reactivity* (the possibility to perceive an environment in which it is situated and respond to perceived changes) and *pro-activeness* (the possibility to take the initiative, starting some activity according to internal goals rather than as a reaction to an external stimulus). This definition, considered by the authors a *weak* definition of agency, is generally already too restrictive to describe as agents most of the entities populating agent-based models for simulation in different fields. Even if the distance between the context of research on *intelligent* agents and agent-based simulation, that is often more focused on the resulting behavior of the local action and interaction of relatively simple agents, cannot be neglected, the aim of this section is to present some relevant results of research on agent models and technologies in computer science and put them in relation with current research on agent-based simulation in other research areas.

As previously suggested, agent-based models (ABMs) can be considered models of complex systems and the ABM approach considers that simple and complex phenomena can be the result of interactions between autonomous and independent entities (i.e., agents) which operate within communities in accordance with different modes of interaction. Thus, agents and ABMs should not be considered simply as a technology (Luck et al. 2005; Zambonelli and Parunak 2003) but also as a modeling approach that can be exploited to represent some system properties that are not simply described as properties or functionalities of single system components but sometimes *emerge from collective behaviors* (Ferber 1999). The study of such emerging behaviors is a typical activity in complex systems modeling (Bar-Yam 1997) and agent-based models are growingly employed for the study of complex systems (see, e.g., Alfi et al. 2007; Hassas et al. 2007; Weyns et al. 2008) satellite workshops of the 2007 edition of the European Conference on Complex Systems).

Agent Behavior Specification

This section will first of all discuss the possible notions of *actions* available to an agent, then it will discuss the way an action actually chooses the

actions to be carried out, introducing the notion of *agent architecture*.

Agent Actions – Actions are the elements at the basis of agent behavior. Agent actions can cause modifications in their environment or in other agents that constitutes the ABM. Different modeling solutions can be provided in order to describe agent actions: as *transformation of a global state*, as *response to influences*, as *computing processes*, as *local modification*, as *physical displacement*, and as *command* (more details about the reported methods to represent agent actions can be found in Ferber (1999)).

- *Functional transformation of states* is based on concepts of states and state transformation and constitutes the most classical approach of artificial intelligence to action representation (mainly used in the planning and multi-agent planning contexts). According to this approach, agent actions are defined as operators whose effect is a change in the state of the world (Fikes and Nilsson 1971; Genesereth and Nilsson 1987; Georgeff 1984).
- Modeling *actions as local modification* provides an approach opposite to the one based on transformation of a global state. Agents perceive their local environment (the only part of the environment that they can access) and according to their perceptions, they modify their internal state. Automata networks (Goles and Martinez 1990) and cellular automata (Wolfram 1986) are examples of this type of models. They are dynamical systems whose behavior is defined in terms of local relationships and local transformations; in turn, these changes influence the overall system state. Within the ABMs context cellular automata are often exploited to represent the dynamic behavior of agent environment (Torrens 2002) or to simulate population dynamics (Epstein and Axtell 1996) in artificial life (Adami 1998; Langton 1995).
- Modeling *actions as response to influences* (Ferber and Muller 1996) extends the previous approach introducing elements to consider the effects of agent interactions and the simultaneous execution of actions in an ABM. Agent actions are conditioned and represent a reaction to other agent actions or to environment modifications.
- Agents can also be considered as *computational processes* (in the vein of the actor model (Agha 1986)). A computing system, and thus a ABM, can be considered as a set of activities (i.e., processes) that are executed sequentially or in parallel. This modeling approach in particular focuses on single computing entities, their behaviors and interactions. The most studied and well known methods to represent processes are finite state automata, Petri nets and their variants (Murata 1989).
- *Approaches derived from physics* can also be found in order to represent agent actions. In these cases, actions mainly concern movements and their applications are in the robotics contexts (i.e., reactive and situated agents) (Latombe 1991). One of the most used notion derived from physics is the one of *field* (e.g., gravitational, electrical, magnetic). Agents are attracted or repelled by given objects or environment areas that emit fields to indicate their position.
- According to cybernetics and control system theory, actions can be represented as *system commands* that regulate and control agent behavior. In this way, actions are complex tasks that the agent executes in order to fulfill given goals and that take into account the environment reactions and correctives to previous actions.

Agent Architecture – The term *architecture* (Russel and Norvig 1995) refers to the model of agent internal structure that is responsible of effectively selecting the actions to be carried out, according to the perceptions and internal state of an agent. Different architectures have been proposed in order to obtain specific agent behaviors and they are generally classified into *deliberative* and *reactive* (respectively, *hysteretic* and *tropistic* according to the classification reported in Genesereth and Nilsson (1987)).

Reactive agents are elementary (and often memory-less) agents with a defined position in the environment. Reactive agents perform their actions as a consequence of the perception of stimuli coming either from other agents or from the environment; generally, the behavioral specification of this kind of agent is a set of condition-action rules, with the addition of a selection strategy for choosing an action to be carried out whenever more rules could be activated. In this case, the motivation for an action derives from a triggering event detected in the environment; these agents cannot be pro-active.

Deliberative or *cognitive* agents, instead, are characterized by a more complex action selection mechanism, and their behavior is based on so called mental states, on facts representing agent knowledge about the environment and, possibly, also on memories of past experiences. Deliberative agents, for every possible sequence of perceptions, try to select a sequence of actions, allowing them to achieve a given goal. Deliberative models, usually defined within the planning context, provide a symbolic and explicit representation of the world within agents and their decisions are based on logic reasoning and symbol manipulation. The BDI model (belief, desire, intention (Rao and Georgeff 1991; Rao and Georgeff 1995)) is perhaps the most widespread model for deliberative agents. The internal state of agents is composed of three “data structures” concerning agent beliefs, desires and intentions. Beliefs represent agent information about its surrounding world, desires are the agent goals, while intentions represent the desire an agent has effectively selected, that it has to some extent committed.

Hybrid architectures can also be defined by combining the previous ones. Agents can have a layered architecture, where deliberative layers are based on a symbolic representation of the surrounding world, generate plans and take decisions, while reactive layers perform actions as effect of the perception of external stimuli. Both vertical and horizontal architectures have been proposed in order to structure layers (Brooks 1986). In horizontal architecture no priorities are associated to layers and the results of the different layers must be combined to obtain agent’s

behavior. When layers are instead arranged in a vertical structure, reactive layers have higher priority over deliberative ones, that are activated only when no reactive behavior is triggered by the perception of an external stimulus.

A MAS can be composed of cognitive agents (generally a relatively low number of deliberative agents), each one possessing its own knowledge-model determining its behavior and its interactions with other agents and the environment. By contrast, there could be MAS made only by reactive agents. This type of system is based on the idea that it is not necessary for a single agent to be individually intelligent for the system to demonstrate complex (intelligent) behaviors. Systems of reactive agents are usually more robust and fault tolerant than other agent-based systems (e.g., an agent may be lost without any catastrophic effect for the system). Other benefits of reactive MAS include flexibility and adaptability in contrast to the inflexibility that sometimes characterizes systems of deliberative agents (Brooks 1990). Finally, a system might also present an *heterogeneous* composition of reactive and deliberative agents.

Environment

Weyns et al. in (Weyns et al. 2007) provide a definition of the notion of environment for MASs (and thus of an environment for an ABM), and also discuss the core responsibilities that can be ascribed to it. In particular, in the specific context of simulation the environment is typically responsible for the following:

- Reflecting/reifying/managing the structure of the physical/social arrangement of the overall system;
- Embedding, supporting regulated access to objects and parts of the system that are not modeled as agents;
- Supporting agent perception and situated action (it must be noted that agent interaction should be considered a particular kind of action);
- Maintain internal dynamics (e.g., spontaneous growth of resources, dissipation signals emitted by agents);
- Define/enforce rules.

In order to exemplify this schema, we will now consider agent-based models and simulators that are based on a physical approach; the latter generally consider agents as particles subject to and generating forces. In this case, the environment comprises laws regulating these influences and relevant elements of the simulated system that are not agents (e.g., point of reference that generate attraction/repulsion forces). It is the environment that determines the overall dynamics, combining the effects that influence each agent and applying them generally in discrete time steps. In this cycle, it captures all the above introduced responsibilities, and the role of agents is minimal (according to some definitions they should not be called agents at all), and running a simulation is essentially reduced to computing iteratively a set equations (see, e.g., (Balmer and Nagel 2006; Helbing et al. 1997)). In situated ABM approaches agents have a higher degree of autonomy and control over their actions, since they evaluate their perceptions and choose their actions according to their behavioral specification. The environment retains a very relevant role, since it provides agents with their perceptions that are generated according to the current structure of the system and to the arrangement of agents situated in it. Socioeconomic models and simulations provide various approaches to the representation of the simulated system, but are generally similar to situated ABMs.


It is now necessary to make a clarification on how the notion of environment in the context of MAS-based simulation can be turned into a software architecture. Klügl et al. (2005) argue that the notion of environment in multi-agent simulation is actually made up of two conceptually different elements: the simulated environment and the simulation environment. The former is a part of the computational model that represents the reality or the abstraction that is the object of the simulation activity. The simulation environment, on the other hand, is a software infrastructure for executing the simulation. In this framework, to make an explicit decoupling between these levels is a prerequisite for good engineering practice. It must be noted that also a different work (Gouaich et al. 2005), non-specifically developed in the

context of agent-based simulation, provided a model for the deployment environment, that is the specific part of the software infrastructure that manages the interactions among agents.

Another recent work is focused on clarifying the notion of ABM environment and describes a three layered model for situated ABM environments (Weyns et al. 2006a). This work argues that environmental abstractions (as well as those related to agents) crosscut all the system levels, from application specific ones, to the execution platform, to the physical infrastructure. There are thus application specific aspects of agents' environment that must be supported by the software infrastructure supporting the execution of the ABM, and in particular the ABM framework (MAS framework in the figure). Figure 3 compares the two above-described schemas.

The fact that the environment actually crosscuts all system levels in a deployment model represents a problem making difficult the separation between simulated environment and simulation infrastructure. In fact, the modeling choices can have a deep influence on the design of the underlying ABM framework and, vice versa, design choices on the simulation infrastructure make it suitable for some ABM and environment models but not usable for other ones. As a result, general ABM framework supporting simulation actually exist, but they cannot offer a specific form of support to the modeler, although they can offer basic mechanisms and abstractions.

SeSAm, for instance, offers a general simulation infrastructure but relies on plugins (Klügl et al. 2005). Those plugins, for example, could be used to define and manage the spatial features of the simulated environment, including the associated basic functions supporting agent movement and perception in that kind of environment. With reference to Fig. 3b, such a plugin would be associated to the application environment module, in the ABM application layer. However, these aspects represent just some of the features of the simulated environment, that can actually comprise rules and laws that extend their influence over the agents and the outcomes of their attempts to act in the environment.



Agent-Based Modeling and Simulation, Fig. 3 A schema introduced in (Klühl et al. 2005) to show differences and relationships between simulated and simulation environment (a), and a three layer deployment model for situated MAS introduced in (Weyns et al. 2006a) highlighting the crosscutting abstractions *agent* and *environment* (b)

Agent Interaction

Interaction is a key aspect in ABM. There is a plethora of definitions for the concept of agent and most of them emphasize the fact that this kind of entity should be able to interact with their environment and with other entities in order to solve problems or simply reach their goals according to coordination, cooperation or competition schemes. The essence of a ABM is the fact that the global system dynamics emerges from the local behaviors and interactions among its composing parts. Strictly speaking, for some kind of ABM the global dynamics is just the sum of local behaviors and interactions, so we cannot always speak of emergent behavior when we talk about a ABM. However the assumptions that underlie the design of an interaction model (or the choice of an existing one for the design and implementation of a specific application) are so important that they have a deep impact on the definition of agent themselves (e.g., an interpreter of a specific language, a perceiver of signals). Therefore it is almost an obvious consequence that interaction mechanisms have a huge impact on the modeling, design and development

of applications based on a specific kind of ABM, which, in turn, is based on a particular interaction model. It is thus not a surprise that a significant part of the research that was carried out in the agent area was focused on this aspect.

This section presents a conceptual taxonomy of currently known/available agent interaction models, trying to define advantages and issues related to them, both from a conceptual and a technical point of view.


There are many possible dimensions and aspects of agent interaction models that can be chosen and adopted in order to define a possible taxonomy. The first aspect that is here considered to classify agent interaction models is related to the fact that agents communicate directly (for instance exchanging messages), and thus the model does not include an abstraction of the actual communication channel, or there are some media interposed among the communication partners which is explicitly included in the interaction model. While the former approach, with specific reference to Agent Communication Language (ACL)-based models, is the most widely adopted in the agent area it has its drawbacks, and

most of them are related to the issue of agents acquaintance. The way ACL-based agent interaction models deal with this issue is the subject of another dimension of the taxonomy, providing direct a priori acquaintance among agents, the adoption of middle agents for information discovery and the development of more complex acquaintance models to tackle issues related to the representation and maintenance of acquaintance information but also to robustness and scalability of the agent infrastructure. However there are other agent interaction models providing an indirect communication among them. Some of these approaches provide the creation and exploitation of artifacts that represent a medium for agents' interaction. Other indirect interaction models are more focused on modeling agent environment as the place where agent interactions take place, thus influencing interactions and agent behavior (Fig. 4).

Direct Interaction Models The first and most widely adopted kind of agent interaction model provide a direct information exchange between communication partners. This approach ignores the details related to the communication channel that allows the interaction, and does not include it as an element of the abstract interaction model. Generally the related mechanisms provide a point-to-point message-passing protocol regulating the exchange of messages among agents. There are various aspects of the communicative act that must be modeled (ranging from low-level technical considerations on message format to conceptual issues related to the formal semantics of messages and conversations), but generally this approach

provides the definition of suitable languages to cover these aspects. While this approach is generally well-understood and can be implemented in a very effective way (especially as it is substantially based on the vast experience of computer networks protocols), in the agent context, it requires specific architectural and conceptual solutions to tackle issues related to the agent acquaintance/discovery and on-tological issues.

Intuitively an Agent Communication Language (ACL) provides agents with a means of exchanging information and knowledge. This vague definition inherently includes the point of view on the conception of the term agent, which assumes that an agent is an intelligent autonomous entity whose features include some sort of *social ability* (Wooldridge and Jennings 1995). According to some approaches this kind of feature is the one that ultimately defines the essence of agency (Genesereth and Ketchpel 1994). Leaving aside the discussion on the definition and conception of agency, this section will focus on what the expression "social ability" effectively means. To do so we will briefly compare these ACL share with those approaches that allow the exchange of information among distributed components (e.g., in legacy systems (With this expression we mean pieces of software which are not designed to interact with agents and agent based systems.)) some basic issues: in particular, the definition of a communication channel allowing the reliable exchange of messages over a computer network (i.e., the lower level aspects of the communication). What distinguishes ACLs from such systems are the objects of




Agent-Based Modeling and Simulation, Fig. 4 The proposed taxonomy of agent interaction models

discourse and their semantic complexity; in particular there are two aspects which distributed computing protocols and architectures do not have to deal with (Fig. 5):

- *Autonomy* of interacting components: modern systems' components (even though they can be quite complex and can be considered as self-sufficient with reference to supplying a specific service) have a lower degree of autonomy than the one that is generally associated to agents.
- The information conveyed in messages does not generally require a comprehensive *ontological approach*, as structures and categories can be considered to be shared by system components.

Regarding the autonomy, while traditional software components offer services and generally perform the required actions as a reaction to the external requests, agents may decide not to carry out a task that was required by some other system entity. Moreover generally agents are considered temporally continuous and proactive, while this is not generally true for common software components.

For what concerns the second point, generally components have specific interfaces which assume an agreement on a set of shared data structures. The semantics of the related information, and the semantics of messages/method invocation/service requests, is generally given on some kind of (more or less formally specified) modeling language, but is tightly related to component implementation. For agent interaction a more explicit and comprehensive view on domain concepts must be specified. In order to be able to effectively exchange knowledge, agents must share an *ontology* (see, e.g., Gruber 1995), that is a representation of a set of categories of objects,




Agent-Based Modeling and Simulation, Fig. 5 Layers and concerns of an agent communication language

concepts, entities, properties and relations among them. In other words, the same concept, object or entity must have a uniform meaning and set of properties across the whole system.

Indirect Interaction Models From a strictly technical point of view, agent communication is generally indirect even in direct agent interaction models. In fact most of these approaches adopt some kind of communication infrastructure supplying a reliable end-to-end message passing mechanism. Nonetheless, the adoption of a *conceptually* direct agent interaction model brings specific issues that were previously introduced. The remaining of this section will focus on models providing the presence of an intermediate entity mediating (allowing and regulating) agent interaction. This communication abstraction is not merely a low-level implementation detail, but a first-class concept of the model (Fig. 6).

Agent interaction models which provide indirect mechanisms of communication will be classified into *artifact mediated* and *spatially grounded* models. The distinction is based on the inspiration and metaphor on which these models are rooted. The former provide the design and implementation of an artifact which emulates concrete objects of agents' environment whose goal is the communication of autonomous entities. Spatially grounded agent interaction models bring the metaphor of modeling agent environment to the extreme, recognizing that there are situations in



Agent-Based Modeling and Simulation, Fig. 6 A conceptual diagram for a typical blackboard architecture, including two sample primitives of the Linda coordination model, that is, the output of a tuple into the blackboard (the *out* operation) and the non-destructive input of an agent from the blackboard (the *read* operation)

which spatial features and information represent a key factor and cannot be neglected in analyzing and modeling a system.

Both of these approaches provide interaction mechanisms that are deeply different from point-to-point message exchange among entities. In fact, the media which enable the interaction intrinsically represent a *context* influencing agent communication.

In the real world a number of physical agents interact sharing resources, by having a competitive access to them (e.g., cars in streets and crossroads), but also collaborating in order to perform tasks which could not be carried out by single entities alone, due to insufficient competencies or abilities (e.g., people that carry a very heavy burden together). Very often, in order to regulate the interactions related to these resources, we build concrete artifacts, such as traffic lights on the streets, or neatly placed handles on large heavy boxes. Exploiting this metaphor, some approaches to agent interaction tend to model and implement abstractions allowing the cooperation of entities through a shared resource, whose access is regulated according to precisely defined rules. *Blackboard-based architectures* are the first examples of this kind of models. A blackboard is a shared data repository that enables cooperating software modules to communicate in an indirect and anonymous way (Englemore and Morgan 1988). In particular the concept of *tuple space*, first introduced in Linda (Gelernter 1985), represents a pervasive modification to the basic blackboard model.

Linda (Gelernter 1985) coordination language probably represents the most relevant blackboard-based model. It is based on the concept of tuple space, that is an associative blackboard allowing agents to share and retrieve data (i.e., tuples) through some data-matching mechanism (such as pattern-matching or unification) integrated within the blackboard. Linda also defines a very simple language defining mechanisms for accessing the tuple space.


The rationale of this approach is to keep separated computation and coordination contexts as much as possible (Gelernter and Carriero 1992), by providing specific abstractions for agent

interaction. With respect to direct interaction models, part of the burden of coordination is in fact moved from the agent to the infrastructure. The evolution of this approach has basically followed two directions: the extension of the coordination language and infrastructure in order to increase its expressiveness or usability, and the modeling and implementation of distributed tuple spaces (Cabri et al. 2000; Omicini and Zambonelli 1999; Picco et al. 1999).

While the previously described indirect approaches define artifacts for agent interaction taking inspiration from actual concrete object of the real world, other approaches bring the metaphor of agent environment to the extreme by taking into account its spatial features.

In these approaches agents are situated in an environment whose spatial features are represented possibly in an explicit way and have an influence on their perception, interaction and thus on their behavior. The concept of perception, which is really abstract and metaphoric in direct interaction models and has little to do with the physical world (agents essentially perceive their state of mind, which includes the effect of received messages, like new facts in their knowledge base), here is related to a more direct modeling of what is often referred to as “local point of view”. In fact these approaches provide the implementation of an infrastructure for agent communication which allows them to perceive the state of the environment in their position (and possibly in nearby locations). They can also cause local modifications to the state of the environment, generally through the emission of signals, emulating some kind of physical phenomenon (e.g., pheromones (Brueckner 2000), or fields (Bandini et al. 2002; Mamei et al. 2002)) or also by simply observing the actions of other agents and reacting to this perception, in a “behavioral implicit communication” schema (Tummolini et al. 2004) (Fig. 7).

In all these cases, however, the structuring function of the environment is central, since it actually defines what can be perceived by an agent in its current position and how it can actually modify the environment, to which extent its actions can be noted by other agents and thus interact with them.



Agent-Based Modeling and Simulation, Fig. 7 A sample schema exemplifying an environment mediated form of interaction in which the spatial structure of the environment has a central role in determining agents' perceptions and their possibility to interact

Platforms for Agent-Based Simulation


Considering the pervasive diffusion and adoption of agent-based approaches to modeling and simulation, it is not surprising the fact that there is a growing interest in software frameworks specifically aimed at supporting the realization of agent-based simulation systems.

This kind of framework often provides abstractions and mechanisms for the definition of agents and their environments, to support their interaction, but also additional functionalities like the management of the simulation (e.g., set-up, configuration, turn management), its visualization, monitoring and the acquisition of data about the simulated dynamics. It is not the aim of this article to provide a detailed review of the current state of the art in this sector, but rather sketch some classes of instruments that have been used to support the realization of agent-based simulations and

provide a set of references to relevant examples of platforms facilitating the development of agent-based simulations.

A first category of these platforms instruments provides *general purpose frameworks* in which agents mainly represent *passive abstractions*, sort of data structures that are manipulated by an overall simulation process. A relevant example of such tools is **NetLogo** (<http://ccl.northwestern.edu/netlogo/>), a dialect of the Logo language specifically aimed at modeling phenomena characterized by a decentralized, interconnected nature. NetLogo does not even adopt the term agent to denote individuals, but it rather calls them *turtles*; a typical simulation consists in a cycle choosing and performing an action for every turtle, considering its current situation and state. It must be noted that, considering some of the previously mentioned definitions of autonomous agent a turtle should not be considered an agent, due to the almost absent autonomy of these entities. The choice of a very simple programming language that does not require a background on informatics, the possibility to deploy in a very simple way simulations as Java applets, and the availability of simple yet effective visualization tools, made NetLogo extremely popular (Fig. 8).

A second category of platforms provides frameworks that are developed with a similar rationale, providing for very similar support tools, but these instruments are based on *general purpose programming languages* (generally object oriented). **Repast** (<http://repast.sourceforge.net/>) (North et al. 2006) represents a successful representative of this category, being a widely employed agent-based simulation platform based on the Java language. The object-oriented nature of the underlying programming language supports the definition of computational elements that make these agents more autonomous, closer to the common definitions of agents, supporting the encapsulation of state (and state manipulation mechanisms), actions and action choice mechanism in agent's class. The choice of adopting a general purpose programming language, on one hand, makes the adoption of these instruments harder for modelers without a background in informatics but, on the other, it



Agent-Based Modeling and Simulation, Fig. 8 A screenshot of a Netlogo simulation applet, (a), and a Repast simulation model, (b)

simplifies the integration with external and existing libraries. Repast, in its current version, can be easily connected to instruments for statistical analysis, data visualization, reporting and also geographic information systems.

While the above-mentioned functionalities are surely important in simplifying the development of an effective simulator, and even if in principle it is possible to adapt frameworks belonging to the previously described categories, it must be noted that their neutrality with respect to the specific adopted agent model leads to a necessary preliminary phase of adaptation of the platform to the specific features of the model that is being implemented. If the latter defines specific abstractions and mechanisms for agents, their decision-making activities, their environment and the way they interact, then the modeler must in general develop proper computational supports to be able to fruitfully employ the platform. These platforms, in fact, are not endowed with specific supports to the realization of agent deliberation mechanisms or infrastructures for interaction models, either direct or indirect (even if it must be noted that all the above platforms generally provide some form of support to agent environment definition, such as grid-like or graph structures).

A third category of platforms represent an attempt to provide a *higher level linguistic support* trying to reduce the distance between agent-based models and their implementations. The latest version of Repast, for instance, is characterized

by the presence of a high level interface for “point-and-click” definition of agent’s behaviors, that is based on a set of primitives for the specification of agent’s actions. **SimSesam** (<http://www.simsesam.de/>) (Klügl et al. 2003) defines a set of *primitive functions* as basic elements for describing agents’ behaviors, and it also provides visual tools supporting model implementation. At the extreme border of this category, we can mention efforts that are specifically aimed at supporting the development of simulations based on a precise agent model, approach and sometimes even for a specific area of application, such as the one described in (Bandini et al. 2007; Weyns et al. 2006b).

Future Directions

Agent-based modeling and simulation is a relatively young yet already widely diffused approach to the analysis, modeling and simulation of complex systems. The heterogeneity of the approaches, modeling styles and applications that legitimately claim to be “agent-based”, as well as the fact that different disciplines are involved in the related research efforts, they are all factors that hindered the definition of a generally recognized view of the field. A higher level framework of this kind of activity would be desirable in order to relate different efforts by means of a shared schema. Moreover it could represent the first step in effectively facing some of the

epistemological issues related to this approach to the modeling and analysis of complex systems. The future directions in this broad research area are thus naturally aimed at obtaining vertical analytical results on specific application domains, but they must also include efforts aimed at “building bridges” between the single disciplinary results in the attempt to reach a more general and shared understanding of how these bottom-up modeling approaches can be effectively employed to study, explain and (maybe) predict the overall behavior of complex systems.

Bibliography

- Adami C (1998) Introduction to artificial life. Springer, New York
- Agha G (1986) Actors: a model of concurrent computation in distributed systems. MIT Press, Cambridge
- Alfi V, Galla T, Marsili M, Pietronero L (eds) (2007) Interacting agents, complexity and interdisciplinary applications (IACIA)
- Balmer M, Nagel K (2006) Shape morphing of intersection layouts using curb side oriented driver simulation. In: van Leeuwen JP, Timmermans HJ (eds) Innovations in design & decision support systems in architecture and urban planning. Springer, Dordrecht, pp 167–183
- Bandini S, Manzoni S, Simone C (2002) Heterogeneous agents situated in heterogeneous spaces. *Appl Artif Intell* 16:831–852
- Bandini S, Manzoni S, Vizzari G (2004) Situated cellular agents: a model to simulate crowding dynamics. *IEICE Trans Inf Syst* E87-D, pp 669–676. Special Issues on Cellular Automata
- Bandini S, Petta P, Vizzari G (eds) (2006a) International symposium on agent based modeling and simulation (ABModSim 2006). Cybernetics and systems. Austrian Society for Cybernetic Studies (2006) 18th European meeting on cybernetics and systems research (EMCSR 2006)
- Bandini S, Celada F, Manzoni S, Puzone R, Vizzari G (2006b) Modelling the immune system with situated agents. In: Apolloni B, Marinaro M, Nicosia G, Tagliaferri R (eds) Proceedings of WIRN/NAIS 2005. Lecture notes in computer science, vol 3931. Springer, Berlin, pp 231–243
- Bandini S, Federici ML, Vizzari G (2007) Situated cellular agents approach to crowd modeling and simulation. *Cybern Syst* 38:729–753
- Bar-Yam Y (1997) Dynamics of complex systems. Addison-Wesley, Reading
- Batty M (2001) Agent based pedestrian modeling. *Env Plan B: Plan Des* 28:321–326
- Bazzan ALC, Wahle J, Klügl F (1999) Agents in traffic modelling – from reactive to social behaviour. In: Burgard W, Christaller T, Cremers AB (eds) KI-99: advances in artificial intelligence, 23rd annual German conference on artificial intelligence, Bonn, 13–15 Sept 1999. Lecture notes in computer science, vol 1701. Springer, Berlin, pp 303–306
- Brooks RA (1986) A robust layered control system for a mobile robot. *IEEE J Robot Autom* 2:14–23
- Brooks RA (1990) Elephants don’t play chess. *Robot Auton Syst* 6:3–15
- Brueckner S (2000) An analytic approach to pheromone-based coordination. In: ICMAS. IEEE Computer Society, pp 369–370
- Cabri G, Leonardi L, Zambonelli F (2000) MARS: a programmable coordination architecture for mobile agents. *IEEE Inter Comp* 4:26–35
- Davidsson P, Logan B, Takadama K (eds) (2005) Multi-agent and multi-agent-based simulation, joint workshop MABS (2004), New York, 19 July 2004. Revised selected papers. In: Davidsson P, Logan B, Takadama K (eds) MABS. Lecture notes in computer science, vol 3415. Springer, Berlin
- Dosi G, Fagiolo G, Roventini A (2006) An evolutionary model of endogenous business cycles. *Comput Econ* 27:3–34
- Edmonds B (2001) The use of models – making MABS more informative. In: Multi-agent-based simulation, second international workshop MABS 2000, Boston, July 2000. Revised and additional papers. Lecture notes in computer science, vol 1979. Springer, Berlin, pp 15–32
- Englemore RS, Morgan T (eds) (1988) Blackboard systems. Addison-Wesley, Reading
- Epstein JM, Axtell R (1996) Growing artificial societies. MIT Press, Boston
- Ferber J (1999) Multi-agent systems. Addison-Wesley, London
- Ferber J, Muller J (1996) Influences and reaction: a model of situated multiagent systems. In: Proceedings of the 2nd international conference on multiagent systems
- Fikes RE, Nilsson NJ (1971) STRIPS: a new approach to the application of theorem proving to problem solving. *Artif Intell* 2:189–208
- Franklin S, Graesser A (1997) Is it an agent, or just a program?: a taxonomy for autonomous agents. In: Müller JP, Wooldridge M, Jennings NR (eds) Intelligent agents III, agent theories, architectures, and languages ECAI’96 workshop (ATAL), Budapest, 12–13 Aug 1996. Lecture notes in computer science, vol 1193. Springer, Berlin, pp 21–36
- Gelernter D (1985) Generative communication in Linda. *ACM Trans Program Lang Syst* 7:80–112
- Gelernter D, Carriero N (1992) Coordination languages and their significance. *Commun ACM* 35:97–107
- Genesereth MR, Ketchpel SP (1994) Software agents. *Commun ACM* 37(7):48–53
- Genesereth MR, Nilsson N (1987) Logical foundations of artificial intelligence. Morgan Kaufmann, San Mateo
- Georgeff M (1984) A theory of action in multi-agent planning. In: Proceedings of the AAAI84, pp 121–125
- Gilbert N, Troitzsch KG (2005) Simulation for the social scientist, 2nd edn. Open University Press, Maidenhead
- Goles E, Martinez S (1990) Neural and automata networks, dynamical behavior and applications. Kluwer, Norwell

- Gouaich A, Michel F, Guiraud Y (2005) MIC: a deployment environment for autonomous agents. In: Environments for multi-agent systems, first international workshop (E4MAS 2004). Lecture notes in computer science, vol 3374. Springer, Berlin, pp 109–126
- Gruber TR (1995) Toward principles for the design of ontologies used for knowledge sharing. *Int J Hum Comput Stud* 43:907–928
- Hales D, Edmonds B, Norling E, Rouchier J (eds) (2003) Multi-agent-based simulation III, 4th international workshop MABS 2003, Melbourne, 14 July 2003. Revised papers. In: Hales D, Edmonds B, Norling E, Rouchier J (eds) MABS. Lecture notes in computer science, vol 2927. Springer, Berlin
- Hassas S, Serugendo GDM, Phan D (eds) (2007) Multi-agents for modelling complex systems (MA4CS). <http://bat710.univ-lyon1.fr/~farmetta/MA4CS07>
- Helbing D, Schweitzer F, Keltsch J, Molnár P (1997) Active walker model for the formation of human and animal trail systems. *Phys Rev E* 56:2527–2539
- Henein CM, White T (2005) Agent-based modelling of forces in crowds. In: Davidsson P, Logan B, Takadama K (eds) Multi-agent and multi-agent-based simulation, joint workshop MABS 2004, New York, 19 July 2004. Revised selected papers. Lecture notes in computer science, vol 3415. Springer, Berlin, pp 173–184
- Klügl F, Herrler R, Oechslein C (2003) From simulated to real environments: how to use sesam for software development. In: Schillo M, Klusch M, Müller JP, Tianfield H (eds) MATES. Lecture notes in computer science, vol 2831. Springer, Berlin, pp 13–24
- Klügl F, Fehler M, Herrler R (2005) About the role of the environment in multi-agent simulations. In: Weyns D, Parunak HVD, Michel F (eds) Environments for multi-agent systems, first international workshop E4MAS 2004, New York, 19 July 2004. Revised selected papers, vol 3374, pp 127–149
- Langton C (1995) Artificial life: an overview. MIT Press, Cambridge
- Latombe JC (1991) Robot motion planning. Kluwer, Boston
- Luck M, McBurney P, Sheory O, Willmott S (eds) (2005) Agent technology: computing as interaction. University of Southampton, Southampton
- Mamei M, Zambonelli F, Leonardi L (2002) Co-fields: towards a unifying approach to the engineering of swarm intelligent systems. In: Engineering societies in the agents world III: third international workshop (ESAW 2002). Lecture notes in artificial intelligence, vol 2577. Springer, Berlin, pp 68–81
- Moss S, Davidsson P (eds) (2001) Multi-agent-based simulation, second international workshop MABS 2000, Boston, July 2000. Revised and additional papers. Lecture notes in computer science, vol 1979. Springer, Berlin
- Murata T (1989) Petri nets: properties, analysis and applications. *Proc IEEE* 77:541–580
- North MJ, Collier NT, Vos JR (2006) Experiences creating three implementations of the repeat agent modeling toolkit. *ACM Trans Model Comp Sim* 16:1–25
- Omicini A, Zambonelli F (1999) Coordination for internet application development. *Autono Agents Multi-Agent Syst* 2:251–269. Special issue: Coordination Mechanisms for Web Agents
- Picco GP, Murphy AL, Roman GC (1999) Lime: Linda meets mobility. In: Proceedings of the 21st international conference on software engineering (ICSE 99). ACM Press, Los Angeles, pp 368–377
- Pyka A, Fagiolo G (2007) Agent-based modelling: a methodology for neo-Schumpeterian economics. In: Hanusch H, Pyka A (eds) Elgar companion to neo-schumpeterian economics. Edward Elgar Publishing, pp 467–487
- Rao A, Georgeff M (1991) Modeling rational agents within a BDI-architecture. In: Proceedings of the knowledge representation and reasoning (KR&R 1991)
- Rao A, Georgeff M (1995) BDI agents: from theory to practice. In: Proceedings of the international conference on multi-agent systems
- Russel S, Norvig P (1995) Artificial intelligence: a modern approach. Prentice Hall, Upper Saddle River
- Schadschneider A, Kirchner A, Nishinari K (2002) CA approach to collective phenomena in pedestrian dynamics. In: Bandini S, Chopard B, Tomassini M (eds) Cellular automata, 5th international conference on cellular automata for research and industry ACRI 2002. Lecture notes in computer science, vol 2493. Springer, Berlin, pp 239–248
- Sichman JS, Antunes L (eds) (2006) Multi-agent-based simulation VI, international workshop MABS 2005, Utrecht, The Netherlands, 25 July 2005. Revised and invited papers. In: Sichman JS, Antunes L (eds) MABS. Lecture notes in computer science, vol 3891. Springer, Berlin
- Sichman JS, Conte R, Gilbert N (eds) (1998) Multi-agent systems and agent-based simulation, first international workshop MABS'98, Paris, 4–6 July 1998. Proceedings. In: Sichman JS, Conte R, Gilbert N (eds): MABS. Lecture notes in computer science, vol 1534. Springer, Berlin
- Sichman JS, Bousquet F, Davidsson P (eds) (2003) Multi-agent-based simulation, third international workshop MABS 2002, Bologna, 15–16 July 2002. Revised papers. In: Sichman JS, Bousquet F, Davidsson P (eds) MABS. Lecture notes in computer science, vol 2581. Springer, Berlin
- Torrens P (2002) Cellular automata and multi-agent systems as planning support tools. In: Geertman S, Stillwell J (eds) Planning support systems in practice. Springer, London, pp 205–222
- Tummolini L, Castelfranchi C, Ricci A, Viroli M, Omicini A (2004) “Exhibitionists” and “voyeurs” do it better: a shared environment approach for flexible coordination with tacit messages. In: Weyns D, Parunak HVD, Michel F (eds) 1st international workshop on environments for multi-agent systems (E4MAS 2004), pp 97–111
- Wahle J, Schreckenberg M (2001) A multi-agent system for online simulations based on real-world traffic data. In: Annual Hawaii international conference on system sciences (HICSS-34). IEEE Computer Society, Los Alamitos
- Weyns D, Holvoet T (2006) From reactive robots to situated multi-agent systems: a historical perspective on the

- role of environment in multi-agent systems. In: Dikenelli O, Gleizes MP, Ricci A (eds) *Engineering societies in the agents world VI*, 6th international workshop ESAW (2005). Lecture notes in computer science, vol 3963. Springer, Berlin, pp 63–88
- Weyns D, Schelfhout K, Holvoet T, Lefever T (2005) Decentralized control of E'GV transportation systems. In: *AAMAS industrial applications*. ACM Press, Utrecht, pp 67–74
- Weyns D, Vizzari G, Holvoet T (2006a) Environments for situated multi-agent systems: beyond infrastructure. In: Weyns D, Parunak HVD, Michel F (eds) *Environments for multi-agent systems II*, second international workshop E4MAS (2005), Utrecht, 25 July 2005. Selected revised and invited papers. Lecture notes in computer science, vol 3830. Springer, Berlin, pp 1–17
- Weyns D, Boucké N, Holvoet T (2006b) Gradient field-based task assignment in an AGV transportation system. In: *AAMAS'06: proceedings of the fifth international joint conference on autonomous agents and multiagent systems*. ACM Press, Hakodate, pp 842–849
- Weyns D, Omicini A, Odell J (2007) Environment as a first class abstraction in multiagent systems. *Auton Agents Multi-Agent Syst* 14:5–30
- Weyns D, Brueckner SA, Demazeau Y (eds) (2008) *Engineering environment-mediated multi-agent systems: international workshop, EEMMAS 2007, Dresden, Oct 2007*. Selected revised and invited papers. Lecture notes in computer science, vol 5049. Springer, Berlin
- Wolfram S (1986) *Theory and applications of cellular automata*. World Press, Singapore
- Wooldridge MJ, Jennings NR (1995) Intelligent agents: theory and practice. *Knowl Eng Rev* 10:115–152
- Zambonelli F, Parunak HVD (2003) Signs of a revolution in computer science and software engineering. In: Petta P, Tolks-dorf R, Zambonelli F (eds) *Engineering societies in the agents world III*, third international workshop, ESAW 2002, Madrid, Sept 2002. Revised papers. Lecture notes in computer science, vol 2577. Springer, Berlin, pp 13–28

Agent-Based Modeling, Mathematical Formalism for

Reinhard Laubenbacher¹, Abdul S. Jarrah²,
Henning S. Mortveit¹ and S. S. Ravi³

¹Virginia Bioinformatics Institute, Virginia
Polytechnic Institute and State University,
Virginia, USA

²Department of Mathematics and Statistics,
American University of Sharjah, Sharjah, United
Arab Emirates

³Department of Computer Science, University at
Albany – State University of New York,
New York, USA

Article Outline

Glossary

Definition of the Subject

Introduction

Existing Mathematical Frameworks

Finite Dynamical Systems

Finite Dynamical Systems as Theoretical and
Computational Tools

Mathematical Results on Finite Dynamical
Systems

Future Directions

Bibliography

Glossary

Agent-based simulation An agent-based simulation of a complex system is a computer model that consists of a collection of agents/variables that can take on a typically finite collection of states. The state of an agent at a given point in time is determined through a collection of rules that describe the agent's interaction with other agents. These rules may be deterministic or stochastic. The agent's state depends on the agent's previous state and the state of a collection of other agents with whom it interacts.

Finite dynamical system A finite dynamical system is a time-discrete dynamical system on a finite state set. That is, it is a mapping from a Cartesian product of finitely many copies of a finite set to itself. This finite set is often considered to be a field. The dynamics is generated by iteration of the mapping.

Mathematical framework A mathematical framework for agent-based simulation consists of a collection of mathematical objects that are considered mathematical abstractions of agent-based simulations. This collection of objects should be general enough to capture the key features of most simulations, yet specific enough to allow the development of a mathematical theory with meaningful results and algorithms.

Definition of the Subject

Agent-based simulations are generative or computational approaches used for analyzing “complex systems.” What is a “system?” Examples of systems include a collection of molecules in a container, the population in an urban area, and the brokers in a stock market. The entities or *agents* in these three systems would be molecules, individuals, and stock brokers, respectively. The agents in such systems interact in the sense that molecules collide, individuals come into contact with other individuals, and brokers trade shares. Such systems, often called *multiagent systems*, are not necessarily complex. The label “complex” is typically attached to a system if the number of agents is large, if the agent interactions are involved, or if there is a large degree of heterogeneity in agent character or their interactions.

This is of course not an attempt to define a complex system. Currently, there is no generally agreed upon definition of complex systems. It is not the goal of this entry to provide such a definition – for our purposes, it will be sufficient to think of a complex system as a collection of agents

interacting in some manner that involves one or more of the complexity components just mentioned, that is, with a large number of agents, heterogeneity in agent character and interactions, and possibly stochastic aspects to all these parts. The global properties of complex systems, such as their global dynamics, emerge from the totality of local interactions between individual agents over time. While these local interactions are well understood in many cases, little is known about the emerging global behavior arising through interaction. Thus, it is typically difficult to construct global mathematical models such as systems of ordinary or partial differential equations, whose properties one could then analyze. Agent-based simulations are one way to create computational models of complex systems that take their place.

An *agent-based simulation*, sometimes also called an *individual-based* or *interaction-based simulation* (which we prefer), of a complex system is in essence a computer program that realizes some (possibly approximate) model of the system to be studied, incorporating the agents and their rules of interaction. The simulation might be deterministic (i.e., the evolution of agent states is governed by deterministic rules) or stochastic. The typical way in which such simulations are used is to initialize the computer program with a particular assignment of agent states and to run it for some time. The output is a temporal sequence of states for all agents, which is then used to draw conclusions about the complex system one is trying to understand. In other words, the computer program is the model of the complex system, and by running the program repeatedly, one expects to obtain an understanding of the characteristics of the complex system.

There are two main drawbacks to this approach. First, it is difficult to validate the model. Simulations for most systems involve quite complex software constructs that pose challenges to code validation. Second, there are essentially no rigorous tools available for an analysis of model properties and dynamics. There is also no widely applicable formalism for the comparison of models. For instance, if one agent-based simulation is a simplification of another, then one

would like to be able to relate their dynamics in a rigorous fashion. We are currently lacking a mathematically rich formal framework that models agent-based simulations. This framework should have at its core a class of mathematical objects to which one can map agent-based simulations. The objects should have a sufficiently general mathematical structure to capture key features of agent-based simulations and, at the same time, should be rich enough to allow the derivation of substantial mathematical results. This entry presents one such framework, namely, the class of *time-discrete dynamical systems over finite state sets*.

The building blocks of these systems consist of a collection of variables (mapping to agents), a graph that captures the dependence relations of agents on other agents, a local update function for each agent that encapsulates the rules by which the state of each agent evolves over time, and an update discipline for the variables (e.g., parallel or sequential). We will show that this class of mathematical objects is appropriate for the representation of agent-based simulations and, therefore, complex systems, and is rich enough to pose and answer relevant mathematical questions. This class is sufficiently rich to be of mathematical interest in its own right and much work remains to be done in studying it. We also remark that many other frameworks such as probabilistic Boolean networks (Shmulevich et al. 2002a) fit inside the framework described here.

Introduction

Computer simulations have become an integral part of today's research and analysis methodologies. The ever-increasing demands arising from the complexity and sheer size of the phenomena studied constantly push computational boundaries, challenge existing computational methodologies, and motivate the development of new theories to improve our understanding of the potential and limitations of computer simulation. Interaction-based simulations are being used to simulate a variety of biological systems such as ecosystems and the immune system, social

systems such as urban populations and markets, and infrastructure systems such as communication networks and power grids.

To model or describe a given system, one typically has several choices in the construction and design of agent-based models and representations. When agents are chosen to be simple, the simulation may not capture the behavior of the real system. On the other hand, the use of highly sophisticated agents can quickly lead to complex behavior and dynamics. Also, use of sophisticated agents may lead to a system that *scales poorly*. That is, a linear increase in the number of agents in the system may require a nonlinear (e.g., quadratic, cubic, or exponential) increase in the computational resources needed for the simulation.

Two common methods, namely, *discrete-event simulation* and *time-stepped simulation*, are often used to implement agent-based models (Bagrodia 1998; Jefferson 1985; Nance 1993). In the discrete-event simulation method, each event that occurs in the system is assigned a time of occurrence. The collection of events is kept in increasing order of their occurrence times. (Note that an event occurring at a certain time may give rise to events which occur later.) When all the events that occur at a particular time instant have been carried out, the simulation clock is advanced to the next time instant in the order. Thus, the differences between successive values of the simulation clock may not be uniform. Discrete-event simulation is typically used in contexts such as queuing systems (Misra 1986). In the time-stepped method of simulation, the simulation clock is always advanced by the same amount. For each value of the simulation clock, the states of the system components are computed using equations that model the system. This method of simulation is commonly used for studying, e.g., fluid flows or chemical reactions. The choice of model (discrete event vs. time stepped) is typically guided by an analysis of the computational speed they can offer, which in turn depends on the nature of the system being modeled (see, e.g., Guo et al. 2000).

Toolkits for general purpose agent-based simulations include Swarm (Ebeling and Schweitzer 2001; Minar et al. 1996) and Repast (North et al.

2006). Such toolkits allow one to specify more complex agents and interactions than would be possible using, e.g., ordinary differential equations models. In general, it is difficult to develop a software package that is capable of supporting the simulation of a wide variety of physical, biological, and social systems.

Standard or classical approaches to modeling are often based on continuous techniques and frameworks such as ordinary differential equations (ODEs) and partial differential equations (PDEs). For example, there are PDE-based models for studying traffic flow (Gupta and Katiyar 2005; Keyfitz 2004; Whitham 1999). These can accurately model the emergence of traffic jams for simple road/intersection configurations through, for example, the formation of shocks. However, these models fail to scale to the size and the specifications required to accurately represent large urban areas. Even if they hypothetically were to scale to the required size, the answers they provide (e.g., car density on a road as a function of position and time) cannot answer questions pertaining to specific travelers or cars. Questions of this type can be naturally described and answered through agent-based models. An example of such a system is TRANSIMS (see section “[TRANSIMS \(Transportation Analysis and Simulation System\)](#)”), where an agent-based simulation scheme is implemented through a cellular automaton model. Another well-known example of the change in modeling paradigms from continuous to discrete is given by lattice gas automata (Frish et al. 1986) in the context of fluid dynamics. Stochastic elements are inherent in many systems, and this usually is reflected in the resulting models used to describe them. A stochastic framework is a natural approach in the modeling of, for example, noise over a channel in a simulation of telecommunication networks (Barrett et al. 2002). In an economic market or a game theoretic setting with competing players, a player may sometimes decide to provide incorrect information. The state of such a player may therefore be viewed and modeled by a random variable. A player may make certain probabilistic assumptions about other players’ environment. In biological systems, certain

features and properties may only be known up to the level of probability distributions. It is only natural to incorporate this stochasticity into models of such systems.

Since applications of stochastic discrete models are common, it is desirable to obtain a better understanding of these simulations both from an application point of view (reliability, validation) and from a mathematical point of view. However, an important disadvantage of agent-based models is that there are few mathematical tools available at this time for the analysis of their dynamics.

Examples of Agent-Based Simulations

In order to provide the reader with some concrete examples that can also be used later on to illustrate theoretical concepts, we describe here three examples of agent-based descriptions of complex systems, ranging from traffic networks to the immune system and voting schemes.

TRANSIMS (Transportation, Analysis and Simulation System)

TRANSIMS is a large-scale computer simulation of traffic on a road network (Nagel and Wagner 2006; Nagel et al. 1997; Rickert et al. 1996). The simulation works at the resolution level of individual travelers and has been used to study large US metropolitan areas such as Portland, OR, Washington, DC, and Dallas/Fort Worth. A TRANSIMS-based analysis of an urban area requires (i) a population, (ii) a location-based activity plan for each individual for the duration of the analysis period, and (iii) a network representation of all transportation pathways of the given area. The data required for (i) and (ii) are generated based on, e.g., extensive surveys and other information sources. The network representation is typically very close to a complete description of the real transportation network of the given urban area.

TRANSIMS consists of two main modules: the *router* and the cellular automaton-based *micro-simulator*. The router maps each activity plan for each individual (obtained typically from activity surveys) into a travel route. The micro-simulator executes the travel routes and sends each

individual through the transportation network so that its activity plan is carried out. This is done in such a way that all constraints imposed on individuals from traffic driving rules, road signals, fellow travelers, and public transportation schedules are respected. The time scale is typically 1 s.

The micro-simulator is the part of TRANSIMS responsible for the detailed traffic dynamics. Its implementation is based on *cellular automata* which are described in more detail in section “Cellular Automata.” Here, for simplicity, we focus on the situation where each individual travels by car. The *road network representation* is in terms of *links* (e.g., road segments) and *nodes* (e.g., intersections). The network description is turned into a cell-network description by discretizing each lane of every link into cells. A cell corresponds to a 7.5 m lane segment and can have up to four neighbor cells (front, back, left, and right).

The *vehicle dynamics* is specified as follows. Vehicles travel with discrete velocities 0, 1, 2, 3, 4, or 5 which are constant between time steps. Each update time step brings the simulation one time unit forward. If the time unit is 1 s, then the maximum speed of $v_{\max} = 5$ cells per time unit corresponds to an actual speed of $5 \times 7.5 \text{ m/s} = 37.5 \text{ m/s}$ which is 135 km/h or approximately 83.9 mph.

Ignoring intersection dynamics, the micro-simulator executes three functions for each vehicle in every update: (a) lane-changing, (b) acceleration, and (c) movement. These functions can be implemented through four cellular automata, one each for lane change decision and execution, one for acceleration, and one for movement. For instance, the acceleration automaton works as follows. A vehicle in TRANSIMS can increase its speed by at most 1 cell per second, but if the road ahead is blocked, the vehicle can come to a complete stop in the same time. The function that is applied to each cell that has a car in it uses the gap ahead and the maximal speed to determine if the car will increase or decrease its velocity. Additionally, a car may have its velocity decreased one unit as determined by a certain deceleration probability. The random deceleration is an important element of producing realistic

traffic flow. A major advantage of this representation is that it leads to very lightweight agents, a feature that is critical for achieving efficient scaling.

C-ImmSim

Next, we discuss an interaction-based simulation that models certain aspects of the human immune system. Comprised of a large number of interacting cells whose motion is constrained by the body's anatomy, the immune system lends itself very well to agent-based simulation. In particular, these models can take into account three-dimensional anatomical variations as well as small-scale variability in cell distributions. For instance, while the number of T cells in the human body is astronomical, the number of antigen-specific T cells, for a specific antigen, can be quite small, thereby creating many spatial inhomogeneities. Also, little is known about the global structure of the system to be modeled.

The first discrete model to incorporate a useful level of complexity was *ImmSim* (Celada and Seiden 1992a, b), developed by Seiden and Celada as a stochastic cellular automaton. The system includes B cells, T cells, antigen-presenting cells (APCs), antibodies, antigens, and antibody-antigen complexes. Receptors on cells are represented by bit strings, and antibodies use bit strings to represent their epitopes and peptides. Specificity and affinity are defined by using bit string similarity. The bit string approach was initially introduced in Farmer et al. (1986). The model is implemented on a regular two-dimensional grid, which can be thought of as a slice of a lymph node, for instance. It has been used to study various phenomena, including the optimal number of human leukocyte antigens in human beings (Celada and Seiden 1992a), the autoimmunity and T lymphocyte selection in the thymus (Morpurgo et al. 1995), antibody selection and hyper-mutation (Celada and Seiden 1996), and the dependence of the selection and maturation of the immune response on the antigen-to-receptor's affinity (Bernaschi et al. 2000). The computational limitations of the Seiden-Celada model have been overcome by a modified model, *C-ImmSim* (Castiglione et al. 1997),

implemented on a parallel architecture. Its complexity is several orders of magnitude larger than that of its predecessor. It has been used to model hypersensitivity to chemotherapy (Castiglione and Agur 2003) and the selection of escape mutants from immune recognition during HIV infection (Bernaschi and Castiglione 2002). In Castiglione et al. (2007), the *C-ImmSim* framework was applied to the study of mechanisms that contribute to the persistence of infection with the Epstein-Barr virus.

A Voting Game

The following example describes a hypothetical voting scheme. The voting system is constructed from a collection of voters. For simplicity, it is assumed that only two candidates, represented by 0 and 1, contest in the election. There are N voters represented by the set $\{v_1, v_2, \dots, v_N\}$. Each voter has a candidate preference or a *state*. We denote the state of voter v_i by x_i . Moreover, each voter knows the preferences or states of some of his or her friends (fellow voters). This friendship relation is captured by the *dependency graph* which we describe later in section “[Definitions, Background, and Examples](#).” Informally, the dependency graph has as vertices the voters with an edge between each pair of voters that are friends.

Starting from an initial configuration of preferences, the voters cast their votes in some order. The candidate that receives the most votes is the winner. A number of rules can be formulated to decide how each voter chooses a candidate. We will provide examples of such rules later, and as will be seen, the outcome of the election is governed by the order in which the votes are cast as well as the structure of the dependency graph.

Existing Mathematical Frameworks

The field of agent-based simulation currently places heavy emphasis on implementation and computation rather than on the derivation of formal results. Computation is no doubt a very useful way to discover potentially interesting behavior and phenomena. However, unless the simulation has been set up very carefully, its outcome does

not formally validate or guarantee the observed phenomenon. It could simply be caused by an artifact of the system model, an implementation error, or some other uncertainty.

A first step in a theory of agent-based simulation is the introduction of a formal framework that, on the one hand, is precise and computationally powerful and, on the other hand, is natural in the sense that it can be used to effectively describe large classes of both deterministic and stochastic systems. Apart from providing a common basis and a language for describing the model using a sound formalism, such a framework has many advantages. At a first level, it helps to clarify the key structure in a system. Domain-specific knowledge is crucial to deriving good models of complex systems, but domain specificity is often confounded by domain conventions and terminology that eventually obfuscate the real structure.

A formal, context independent framework also makes it easier to take advantage of existing general theory and algorithms. Having a model formulated in such a framework also makes it easier to establish results. Additionally, expressing the model using a general framework is more likely to produce results that are widely applicable. This type of framework also supports implementation and validation. Modeling languages like UML (Booch et al. 2005) serve a similar purpose but tend to focus solely on software implementation and validation issues and very little on mathematical or computational analysis.

Cellular Automata

In this section, we discuss several existing frameworks for describing agent-based simulations. Cellular automata (CA) were introduced by Ulam and von Neumann (von Neumann and Burks 1966) as biologically motivated models of computation. Early research addressed questions about the computational power of these devices. Since then, their properties have been studied in the context of dynamical systems (Hedlund 1969), language theory (Lindgren et al. 1998), and ergodic theory (Lind 1984), to mention just a few areas. Cellular automata were popularized by Conway (Gardner 1970) (Game of Life) and by Wolfram (Martin et al. 1984; Wolfram 1983, 2002). Cellular automata (both deterministic and stochastic) have been

used as models for phenomena ranging from lattice gases (Frish et al. 1986) and flows in porous media (Rothman 1988) to traffic analysis (Fukš 2004; Nagel and Schreckenberg 1992; Nagel et al. 1995).

A cellular automaton is typically defined over a regular grid. An example is a two-dimensional grid such as $\mathbb{Z}^2$. Each grid point (i, j) is referred to as a *site* or *node*. Each site has a state $x_{i,j}(t)$ which is often taken to be binary. Here, t denotes the time step. Furthermore, there is a notion of a *neighborhood* for each site. The neighborhood N of a site is the collection of sites that can influence the future state of the given site. Based on its current state $x_{i,j}(t)$ and the current states of the sites in its neighborhood N , a function $f_{i,j}$ is used to compute the next state $x_{i,j}(t+1)$ of the site (i, j) . Specifically, we have

$$x_{i,j}(t+1) = f_{i,j}(\bar{x}_{i,j}(t)), \quad (1)$$

where $\bar{x}_{i,j}(t)$ denotes the tuple consisting of all the states $x_{i',j'}(t)$ with $(i', j') \in N$. The tuple consisting of the states of all the sites is the CA *configuration* and is denoted $x(t) = (x_{i,j}(t))_{i,j}$. Equation 1 is used to map the configuration $x(t)$ to $x(t+1)$. The cellular automaton map or dynamical system is the map Φ that sends $x(t)$ to $x(t+1)$.

A central part of CA research is to understand how configurations evolve under iteration of the map Φ and what types of dynamical behavior can be generated. A general introduction to CA can be found in Ilachinsky (2001).

Hopfield Networks

Hopfield networks were proposed as a simple model of associative memories (Hopfield 1982). A discrete Hopfield neural network consists of an undirected graph $Y(V, E)$. At any time t , each node $v_i \in V$ has a state $x_i(t) \in \{+1, -1\}$. Further, each node $v_i \in V$ has an associated *threshold* $\tau_i \in \mathbf{R}$. Each edge $\{v_i, v_j\} \in E$ has an associated weight $w_{i,j} \in \mathbf{R}$. For each node v_i , the neighborhood N_i of v_i includes v_i and the set of nodes that are adjacent to v_i in Y . Formally,

$$N_i = \{v_i\} \cup \{v_j \in V : \{v_i, v_j\} \in E\}.$$

States of nodes are updated as follows. At time t , node v_i computes the function f_i defined by

$$f_i(t) = \text{sgn} \left(-\tau_i + \sum_{v_j \in N_i} w_{ij} x_j(t) \right),$$

where sgn is the map from $\mathbf{R}$ to $\{+1, -1\}$, defined by

$$\text{sgn}(x) = \begin{cases} 1, & \text{if } x \geq 0 \\ -1, & \text{otherwise.} \end{cases}$$

Now, the state of v_i at time $t + 1$ is

$$x_i(t + 1) = f_i(t).$$

Many references on Hopfield networks (see, e.g., Hopfield 1982; Russell and Norwig 2003) assume that the underlying undirected graph is complete; that is, there is an edge between pairs of nodes. In the definition presented above, the graph need not be complete. However, this does not cause any difficulties since the missing edges can be assigned weight 0. As a consequence, such edges will not play any role in determining the dynamics of the system. Both synchronous and asynchronous update models of Hopfield neural networks have been considered in the literature. For theoretical results concerning Hopfield networks, see Orponen (1994, 1996) and the references cited therein. Reference Russell and Norwig (2003) presents a number of applications of neural networks. In Macy et al. (2003), a Hopfield model is used to study polarization in dynamic networks.

Communicating Finite-State Machines

The model of communicating finite-state machines (CFSM) was proposed to analyze protocols used in computer networks. In some of the literature, this model is also referred to as “concurrent transition systems” (Gouda and Chang 1986).

In the CFSM model, each agent is a process executing on some node of a distributed computing system. Although there are minor differences among the various CFSM models proposed in the literature (Brand and Zafropulo 1983; Gouda and Chang 1986), the basic setup models each process as a finite-state machine (FSM). Thus, each agent is in a certain state at any time instant t . For each

pair of agents, there is a bidirectional channel through which they can communicate. The state of an agent at time $t + 1$ is a function of the current state and the input (if any) on one or more of the channels. When an agent (FSM) undergoes a transition from one state to another, it may also choose to send a message to another agent or receive a message from an agent. In general, such systems can be synchronous or asynchronous. As can be seen, CFSMs are a natural formalism for studying protocols used in computer networks. The CFSM model has been used extensively to prove properties (e.g., deadlock freedom, fairness) of a number of protocols used in practice (see Brand and Zafropulo 1983; Gouda and Chang 1986 and the references cited therein).

Other frameworks include interacting particle systems (Liggett 2005) and Petri nets (Moncion et al. 2006). There is a vast literature on both, but space limitations prevent a discussion here.

Finite Dynamical Systems

Another, quite general, modeling framework that has been proposed is that of *finite dynamical systems*, both synchronous and asynchronous. Here, the proposed mathematical object representing an agent-based simulation is a time-discrete dynamical system on a finite state set. The description of the systems is modeled after the key components of an agent-based simulation, namely, agents, the dependency graph, local update functions, and an update order. This makes a mapping to agent-based simulations natural. In the remainder of this entry, we will show that finite dynamical systems satisfy our criteria for a good mathematical framework in that they are general enough to serve as a broad computing tool and mathematically rich enough to allow the derivation of formal results.

Definitions, Background, and Examples

Let $x_1, \dots, x_n$ be a collection of variables, which take values in a finite set X . (As will be seen, the variables represent the entities in the system being modeled and the elements of X represent their states.) Each variable x_i has associated to it a “local update function” $f_i : X^n \rightarrow X$, where “local” refers to the fact that f_i takes inputs from

the variables in the “neighborhood” of x_i , in a sense to be made precise below. By abuse of notation, we also let f_i denote the function $X^n \rightarrow X$ which changes the i th coordinate and leaves the other coordinates unchanged. This allows for the sequential composition of the local update functions. These functions assemble to a dynamical system

$$\Phi = (f_1, \dots, f_n) : X^n \rightarrow X^n,$$

with the dynamics generated by iteration of Φ . As an example, if $X = \{0, 1\}$ with the standard Boolean operators AND and OR, then Φ is a Boolean network.

The assembly of Φ from the local functions f_i can be done in one of several ways. One can update each of the variables simultaneously, that is,

$$\Phi(x_1, \dots, x_n) = (f_1(x_1, \dots, x_n), \dots, f_n(x_1, \dots, x_n)).$$

In this case, one obtains a *parallel dynamical system*.

Alternatively, one can choose to update the states of the variables according to some fixed update order, for example, a permutation $(\pi_1, \pi_2, \dots, \pi_n)$ of the set $\{1, \dots, n\}$. More generally, one could use a *word* on the set $\{1, \dots, n\}$, that is, $\pi = (\pi_1, \dots, \pi_t)$ where t is the length of the word. The function composition

$$\Phi_\pi = f_{\pi_t} \circ f_{\pi_{t-1}} \circ \dots \circ f_{\pi_1}, \quad (2)$$

is called a *sequential dynamical system (SDS)*, and as before, the dynamics of Φ_π is generated by iteration. The case when π is a permutation on $\{1, \dots, n\}$ has been studied extensively (Barrett and Reidys 1999; Barrett et al. 2000, 2001b, 2003c). It is clear that using a different permutation or word σ may result in a different dynamical system Φ_σ . Using a word rather than a permutation allows one to capture the case where some vertices have states that are updated more frequently than others.

Remark 1 Notice that for a fixed π , the function Φ_π is a parallel dynamical system: once the update order π is chosen and the local update functions are composed according to π , that is, the function Φ_π has been computed, then $\Phi_\pi(x_1, \dots, x_n) = g(x_1, \dots,$

$x_n)$ where g is a parallel update dynamical system. However, the maps g_i are not local functions.

The dynamics of Φ is usually represented as a directed graph on the vertex set X^n , called the *phase space* of Φ . There is a directed edge from $\mathbf{v} \in X^n$ to $\mathbf{w} \in X^n$ if and only if $\Phi(\mathbf{v}) = \mathbf{w}$. A second graph that is usually associated with a finite dynamical system is its *dependency graph* $Y(V, E)$. In general, this is a directed graph, and its vertex set is $V = \{1, \dots, n\}$. There is a directed edge from i to j if and only if x_i appears in the function f_j . In many situations, the interaction relationship between pairs of variables is symmetric; that is, variable x_i appears in f_j if and only if x_j appears in f_i . In such cases, the dependency graph can be thought of as an undirected graph. We recall that the dependency graphs mentioned in the context of the voting game (see section “[A Voting Game](#)”) and Hopfield networks (see section “[Hopfield Networks](#)”) are undirected graphs. The dependency graph plays an important role in the study of finite dynamical systems and is sometimes listed explicitly as part of the specification of Φ .

Example 2 Let $X = \{0, 1\}$ (the Boolean case). Suppose we have four variables and the local Boolean update functions are


$$\begin{aligned} f_1 &= x_1 + x_2 + x_3 + x_4, \\ f_2 &= x_1 + x_2, \\ f_3 &= x_1 + x_3, \\ f_4 &= x_1 + x_4, \end{aligned}$$

where “+” represents sum modulo 2. The dynamics of the function $\Phi = (f_1, \dots, f_4) : X^4 \rightarrow X^4$ is the directed graph in Fig. 1a, while the dependency graph is in Fig. 1b.


Example 3 Consider the local functions in Example 2 and let $\pi = (2, 1, 3, 4)$. Then

$$\Phi_\pi = f_4 \circ f_3 \circ f_1 \circ f_2 : X^4 \rightarrow X^4.$$

The phase space of Φ_π is the directed graph in Fig. 2a, while the phase space of Φ_γ , where $\gamma = id$, is in Fig. 2b.



Agent-Based Modeling, Mathematical Formalism for, Fig. 1 The phase space of the parallel system Φ (a) and dependency graph of the Boolean functions from Example 2 (b)



Agent-Based Modeling, Mathematical Formalism for, Fig. 2 The phase spaces from Example 3: Φ_π (a) and Φ_{id} (b)

Notice that the phase space of any function is a directed graph in which every vertex has out-degree one; this is a characteristic property of deterministic functions.

Making use of Boolean arithmetic is a powerful tool in studying Boolean networks, which is not available in general. In order to have an enhanced set of tools available, it is often natural to make an additional assumption regarding X , namely, that it is a finite number system, a *finite field* (Lidl and Niederreiter 1997). This amounts to the assumption that there are “addition” and “multiplication” operations defined on X that satisfy the same rules as ordinary addition and multiplication of numbers. Examples include $\mathbf{Z}_p$, the integers modulo a prime p . This assumption can be thought of as the discrete analog of imposing a coordinate system on an affine space.

When X is a finite field, it is easy to show that for any local function g , there exists a polynomial h such that $g(x_1, \dots, x_n) = h(x_1, \dots, x_n)$ for all $(x_1, \dots, x_n) \in X^n$. To be precise, suppose X is a finite field with q elements. Then

$$g(x_1, \dots, x_n) = \sum_{(c_1, \dots, c_n) \in X^n} g(c_1, \dots, c_n) \prod_{i=1}^n (1 - (x_i - c_i)^{q-1}). \quad (3)$$

This observation has many useful consequences, since polynomial functions have been studied extensively and many analytical tools are available.

Notice that cellular automata and Boolean networks, both parallel and sequential, are special classes of polynomial dynamical systems. In fact, it is straightforward to see that

$$\begin{aligned} x \wedge y &= x \cdot y & x \vee y &= x + y + xy \\ \text{and } \neg x &= x + 1. \end{aligned} \quad (4)$$

Therefore, the modeling framework of finite dynamical systems includes that of cellular automata discussed earlier. Also, since a Hopfield network is a function $X^n \rightarrow X^n$, which can be represented through its local constituent functions, it follows that Hopfield networks also are special cases of finite dynamical systems.

Stochastic Finite Dynamical Systems

The deterministic framework of finite dynamical systems can be made stochastic in several different ways, making one or more of the system’s defining data stochastic. For example, one could use one or both of the following criteria.

- Assume that each variable has a nonempty set of local functions assigned to it, together with a

probability distribution on this set, and each time a variable is updated, one of these local functions is chosen at random to update its state. We call such systems *probabilistic finite dynamical systems* (PFDS), a generalization of probabilistic Boolean networks (Shmulevich et al. 2002b).

- Fix a subset of permutations $T \subseteq S_n$ together with a probability distribution. When it is time for the system to update its state, a permutation $\pi \in T$ is chosen at random, and the agents are updated sequentially using π . We call such systems *stochastic finite dynamical systems* (SFDS).

Remark 4 By Remark, each system Φ_π is a parallel system. Hence a SFDS is nothing but a set of parallel dynamical systems $\{\Phi_\pi : \pi \in T\}$, together with a probability distribution. When it is time for the system to update its state, a system Φ_π is chosen at random and used for the next iteration.

To describe the phase space of a stochastic finite dynamical system, a general method is as follows. Let Ω be a finite collection of systems $\Phi_1, \dots, \Phi_t$, where $\Phi_i : X^n \rightarrow X^n$ for all i , and consider the probabilities $p_1, \dots, p_t$ which sum to 1. We obtain the stochastic phase space

$$\Gamma_\Omega = p_1\Gamma_1 + p_2\Gamma_2 + \dots + p_t\Gamma_t, \quad (5)$$

where Γ_i is the phase space of Φ_i . The associated probability space is $\mathcal{F} = (\Omega, 2^\Omega, \mu)$, where the probability measure μ is induced by the probabilities p_i . It is clear that the stochastic phase space can


be viewed as a *Markov chain* over the state space X^n . The adjacency matrix of Γ_Ω directly encodes the Markov transition matrix. This is of course not new and has been done in, e.g., (Dawson 1974; Shmulevich et al. 2002b; Vasershtein 1969). But we emphasize the point that even though SFDS give rise to Markov chains, *our study of SFDS is greatly facilitated by the rich additional structure* available in these systems. To understand the effect of structural components such as the topology of the dependency graph or the stochastic nature of the update, it is important to study them not as Markov chains but as SFDS.

Example 5. Consider Φ_π and Φ_γ from Example 3 and let Γ_π and Γ_γ be their phase spaces as shown in Fig. 2. Let $p_1 = p_2 = 1/2$. The phase space $(1/2)\Gamma_\pi + (1/2)\Gamma_\gamma$ of the stochastic sequential dynamical system obtained from Φ_π and Φ_γ (with equal probabilities) is presented in Fig. 3.

Agent-Based Simulations as Finite Dynamical Systems

In the following, we describe the generic structure of the systems typically modeled and studied through agent-based simulations. The central notion is naturally that of an *agent*.

Each agent carries a state that may encode its preferences, internal configuration, perception of its environment, and so on. In the case of TRANSIMS, for instance, the agents are the cells making up the road network. The cell state contains information about whether or not the cell is occupied by a vehicle as well as the velocity of the vehicle. One may assume that each cell takes on states



Agent-Based Modeling, Mathematical Formalism for, Fig. 3 The stochastic phase space for Example 5 induced by the two deterministic phase spaces of Φ_π and Φ_γ from Fig. 2. For simplicity, the weights of the edges have been omitted

from the same set of possible states, which may be chosen to support the structure of a finite field.

The agents interact with each other, but typically an agent only interacts with a small subset of agents, its *neighbors*. Through such an interaction, an agent may decide to change its state based on the states (or other aspects) of the agents with which it interacts. We will refer to the process where an agent modifies its state through interaction as an *agent update*. The precise way in which an agent modifies its state is governed by the nature of the particular agent. In TRANSIMS, the neighbors of a cell are the adjacent road network cells. From this adjacency relation, one obtains a dependency graph of the agents. The local update function for a given agent can be obtained from the rules governing traffic flow between cells.

The updates of all the agents may be scheduled in different ways. Classical approaches include synchronous, asynchronous, and event-driven schemes. The choice will depend on system properties or particular considerations about the simulation implementation.

In the case of *CImmSim*, the situation is somewhat more complicated. Here, the agents are also the spatial units of the system, each representing a small volume of lymph tissue. The total volume is represented as a two-dimensional CA, in which every agent has four neighbors, so that the dependency graph is a regular two-dimensional grid. The state of each agent is a collection of counts for the various immune cells and pathogens that are present in this particular agent (volume). Movement between cells is implemented as diffusion. Immune cells can interact with each other and with pathogens while they reside in the same volume. Thus, the local update function for a given cell of the simulation is made up of the two components of movement between cells and interactions within a cell. For instance, a B cell could interact with the Epstein-Barr virus in a given volume and perform a transition from uninfected to infected by the next time step. Interactions as well as movement are stochastic, resulting in a stochastic finite dynamical system. The update order is parallel.

Example 6 The Voting Game (see section “[A Voting Game](#)”) The following scenario is

constructed to illustrate how implementation choices for the system components have a clear and direct bearing on the dynamics and simulation outcomes.

Let the voter dependency graph be the star graph on 5 vertices with center vertex a and surrounding vertices b , c , d , and e . Furthermore, assume that everybody votes opportunistically using the majority rule: the vote cast by an individual is equal to the preference of the majority of his/her friends with the person's own preference included. For simplicity, assume candidate 1 is preferred in the case of a tie.

If the initial preference is $x_a = 1$ and $x_b = x_c = x_d = x_e = 0$, then if voter a goes first, he/she will vote for candidate 0 since that is the choice of the majority of the neighbor voters. However, if b and c go first, they only know a 's preference. Voter b therefore casts his/her vote for candidate 1 as does c . Note that this is a tie situation with an equal number of preferences for candidate 1 (a) and for candidate 2 (b). If voter a goes next, then the situation has changed: the preference of b and c has already changed to 1. Consequently, voter a picks candidate 1. At the end of the day, candidate 1 is the election winner, and the choice of update order has tipped the election!

This example is of course constructed to illustrate our point. However, in real systems, it can be much more difficult to detect the presence of such sensitivities and their implications. A solid mathematical framework can be very helpful in detecting such effects.

Finite Dynamical Systems as Theoretical and Computational Tools

If finite dynamical systems are to be useful as a modeling paradigm for agent-based simulations, it is necessary that they can serve as a fairly universal model of computation. We discuss here how such dynamical systems can mimic Turing machines (TMs), a standard universal model for computation. For a more thorough exposition, we refer the reader to the series of papers by Barrett et al. (2003a, b, 2006, 2007a, b). To make the discussion reasonably self-contained, we provide

a brief and informal discussion of the TM model. Additional information on TMs can be found in any standard text on the theory of computation (e.g., Sipser 1997).

A Computational View of Finite Dynamical Systems: Definitions

In order to understand the relationship of finite dynamical systems to TMs, it is important to view such systems from a computational stand point. Recall that a finite dynamical system $\Phi : X^n \rightarrow X^n$, where X is a finite set, has an underlying dependency graph $Y(V, E)$. From a computational point of view, the nodes of the dependency graph (the agents in the simulation) are thought of as devices that compute appropriate functions. For simplicity, we will assume in this section that the dependency graph is undirected, that is, all dependency relations are symmetric. At any time, the state value of each node $v_i \in V$ is from the specified domain X . The inputs to f_i are the current state of v_i and the states of the neighbors of v_i as specified by Y . The output of f_i , which is also a member of X , becomes the state of v_i at the next time instant. The discussion in this section will focus on sequentially updated systems (SDS), but all results discussed apply to parallel systems as well.

Each step of the computation carried out by an SDS can be thought as consisting of n “mini steps”; in each mini step, the value of the local transition function at a node is computed and the state of the node is changed to the computed value. Given an SDS Φ , a *configuration* C of Φ is a vector $(c_1, c_2, \dots, c_n) \in X^n$. It can be seen that each computational step of an SDS causes a transition from one configuration to another.

Configuration Reachability Problem for SDSs

Based on the computational view, a number of different problems can be defined for SDSs (see for, e.g., Barrett et al. 2001a, 2006, 2007b). To illustrate how SDSs can model TM computations, we will focus on one such problem, namely, the *configuration reachability* (CR) problem: Given an SDS Φ , an initial configuration C and another configuration C' , will Φ , starting from C , ever reach configuration C' ? The problem can also be expressed in terms of the phase space of Φ . Since

configurations such as C and C' are represented by nodes in the phase space, the CR problem boils down to the question of whether there is a directed path in the phase space from C to C' . This abstract problem can be mapped to several problems in the simulation of multiagent systems. Consider, for example, the TRANSIMS context. Here, the initial configuration C may represent the state of the system early in the day (when the traffic is very light) and C' may represent an “undesirable” state of the system (such as heavy traffic congestion). Similarly, in the context of modeling an infectious disease, C may represent the initial onset of the disease (when only a small number of people are infected), and C' may represent a situation where a large percentage of the population is infected. The purpose of studying computational problems such as CR is to determine whether one can efficiently predict the occurrence of certain events in the system from a description of the system. If computational analysis shows that the system can indeed reach undesirable configurations as it evolves, then one can try to identify steps needed to deal with such situations.

Turing Machines: A Brief Overview

A Turing machine (TM) is a simple and commonly used model for general purpose computational devices. Since our goal is to point out how SDSs can also serve as computational devices, we will present an informal overview of the TM model. Readers interested in a more formal description may consult (Sipser 1997).

A TM consists of a set Q of states, a one-way infinite input tape and a read/write head that can read and modify symbols on the input tape. The input tape is divided into cells and each cell contains a symbol from a special finite alphabet. An input consisting of n symbols is written on the leftmost n cells of the input tape. (The other cells are assumed to contain a special symbol called *blank*.) One of the states in Q , denoted by q_s , is the designated *start* state. Q also includes two other special states, denoted by q_a (the *accepting* state) and q_r (the *rejecting* state). At any time, the machine is in one of the states in Q . The *transition function* for the TM specifies for each combination of the current state and the current symbol

under the head, a new state, a new symbol for the current cell (which is under the head), and a movement (i.e., left or right by one cell) for the head. The machine starts in state q_s with the head on the first cell of the input tape. Each step of the machine is carried out in accordance with the transition function. If the machine ever reaches either the accepting or the rejecting state, it halts with the corresponding decision; otherwise, the machine runs forever.

A *configuration* of a TM consists of its current state, the current tape contents, and the position of the head. Note that the transition function of a TM specifies how a new configuration is obtained from the current configuration.

The above description is for the basic TM model (also called *single-tape* TM model). For convenience in describing some computations, several variants of the above basic model have been proposed. For example, in a *multi-tape* TM, there are one or more *work tapes* in addition to the input tape. The work tapes can be used to store intermediate results. Each work tape has its own read/write head, and the definitions of configuration and transition function can be suitably modified to accommodate work tapes. While such an enhancement to the basic TM model makes it easier to carry out certain computations, it does not add to the machine's computational power. In other words, any computation that can be carried out using the enhanced model can also be carried out using the basic model.

As in the case of dynamical systems, one can define a *configuration reachability* (CR) problem for TMs: Given a TM M , an initial configuration $\mathcal{I}_M$ and another configuration C_M , will the TM starting from $\mathcal{I}_M$ ever reach C_M ? We refer to the CR problem in the context of TMs as CR-TM. In fact, it is this problem for TMs that captures the essence of what can be effectively computed. In particular, by choosing the state component of C_M to be one of the halting states (q_a or q_r), the problem of determining whether a function is computable is transformed into an appropriate CR-TM problem. By imposing appropriate restrictions on the resources used by a TM (e.g., the number of steps, the number of cells on the work tapes), one obtains different versions of the

CR-TM problem which characterize different computational complexity classes (Sipser 1997).

How SDSs Can Mimic Turing Machines

The above discussion points out an important similarity between SDSs and TMs. Under both of these models, each computational step causes a transition from one configuration to another. It is this similarity that allows one to construct a discrete dynamical system Φ that can simulate a TM. Typically, each step of a TM is simulated by a short sequence of successive iterations Φ . As part of the construction, one also identifies a suitable mapping between the configurations of the TM being simulated and those of the dynamical system. This is done in such a way that the answer to CR-TM problem is “yes” if and only if the answer to the CR problem for the dynamical system is also “yes.”

To illustrate the basic ideas, we will informally sketch a construction from (Barrett et al. 2006). For simplicity, this construction produces an SDS Φ that simulates a restricted version of TMs; the restriction being that for any input containing n symbols, the number of work tape cells that the machine may use is bounded by a linear function of n . Such a TM is called a *linear bounded automaton* (LBA) (Sipser 1997). Let M denote the given LBA and let n denote the length of the input to M . The domain X for the SDS Φ is chosen to be a finite set based on the allowed symbols in the input to the TM. The dependency graph is chosen to be a simple path on n nodes, where each node serves as a representative for a cell on the input tape. The initial and final configurations C and C' for Φ are constructed from the corresponding configurations of M . The local transition function for each node of the SDS is constructed from the given transition function for M in such a way that each step of M corresponds to exactly one step of Φ . Thus, there is a simple bijection between the sequence of configurations that M goes through during its computation and the sequence of states that Φ goes through as it evolves. Using this bijection, it is shown in Barrett et al. (2006) that the answer to the CR-TM problem is “yes” if and only if Φ reaches C' starting from C . Reference (Barrett et al. 2006) also

presents a number of sophisticated constructions where the resulting dynamical system is very simple; for example, it is shown that an LBA can be simulated by an SDS in which X is the Boolean field, the dependency graph is d -regular for some constant d , and the local transition functions at all the nodes are *identical*. Such results point out that one does not need complicated dynamical systems to model TM computations.

References (Barrett et al. 2003a, b, 2006) present constructions that show how more general models of TMs can also be simulated by appropriate SDSs. As one would expect, these constructions lead to SDSs with more complex local transition functions.

Barrett et al. (2007a) have also considered *stochastic* SDS (SSDS), where the local transition function at each node is stochastic. For each combination of inputs, a stochastic local transition function specifies a probability distribution over the domain of state values. It is shown in Barrett et al. (2007a) that SSDSs can effectively simulate computations carried out by probabilistic TMs (i.e., TMs whose transition functions are stochastic).

TRANSIMS-Related Questions

Section “TRANSIMS (Transportation, Analysis, Simulation System)” gave an overview of some aspects of the TRANSIMS model. The micro-simulator module is specified as a functional composition of four cellular automata of the form $\Delta_4 \circ \Delta_3 \circ \Delta_2 \circ \Delta_1$. (We only described Δ_3 which corresponds to velocity updates.) Such a formulation has several advantages. First, it has created an abstraction of the essence of the system in a precise mathematical way without burying it in contextual domain details. An immediate advantage of this specification is that it makes the whole implementation process more straightforward and transparent. While the local update functions for the cells are typically quite simple, for any realistic study of an urban area, the problem size would typically require a sequential implementation, raising a number of issues that are best addressed within a mathematical framework like the one considered here.

Mathematical Results on Finite Dynamical Systems

In this section, we outline a collection of mathematical results about finite dynamical systems that is representative of the available knowledge. The majority of these results are about deterministic systems, as the theory of stochastic systems of this type is still in its infancy. We will first consider synchronously updated systems.

Throughout this section, we make the assumption that the state set X carries the algebraic structure of a finite field. Accordingly, we use the notation k instead of X . It is a standard result that in this case the number q of elements in k has the form $q = p^t$ for some prime p and $t \geq 1$. The reader may keep the simplest case $k = \{0, 1\}$ in mind, in which case we are effectively considering Boolean networks.

Recall Eq. 4. That is, any function $g: k^n \rightarrow k$ can be represented by a multivariate polynomial with coefficients in k . If we require that the exponent of each variable be less than q , then this representation is unique. In particular, Eq. 4 implies that every Boolean function can be represented uniquely as a polynomial function.

Parallel Update Systems

Certain classes of finite dynamical systems have been studied extensively, in particular cellular automata and Boolean networks where the state set is $\{0, 1\}$. Many of these studies have been experimental in nature, however. Some more general mathematical results about cellular automata can be found in the papers of Wolfram and collaborators (Wolfram 1986). The results there focus primarily on one-dimensional Boolean cellular automata with a particular fixed initial state. Here, we collect a sampling of more general results.

We first consider linear and affine systems

$$\Phi = (f_1, \dots, f_n) : k^n \rightarrow k^n.$$

That is, we consider systems for which the coordinate functions f_i are linear, resp. affine, polynomials. (In the Boolean case, this includes functions constructed from XOR (sum modulo 2)

and negation.) When each f_i is a linear polynomial of the form $f_i(x_1, \dots, x_n) = a_{i1}x_1 + \dots + a_{in}x_n$, the map Φ is nothing but a *linear transformation* on k^n over k , and, by using the standard basis, Φ has the matrix representation

$$\Phi \left(\begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \right) = \begin{bmatrix} a_{11} & \dots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \dots & a_{nn} \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix},$$

where $a_{ij} \in k$ for all i, j .

Linear finite dynamical systems were first studied by Elspas (1959). His motivation came from studying feedback shift register networks and their applications to radar and communication systems and automatic error correction circuits. For linear systems over finite fields of prime cardinality, that is, q is a prime number, Elspas showed that the exact number and length of each limit cycle can be determined from the *elementary divisors* of the matrix A . Recently, Hernández-Toledo (2005) rediscovered Elspas' results and generalized them to arbitrary finite fields. Furthermore, he showed that the structure of the *tree of transients* at each node of each limit cycle is the same and can be completely determined from the nilpotent elementary divisors of the form x^a . For *affine* Boolean networks (i.e., finite dynamical systems over the Boolean field with two elements, whose local functions are linear polynomials which might have constant terms), a method to analyze their cycle length has been developed in Milligan and Wilson (1993). After embedding the matrix of the transition function, which is of dimension $n \times (n + 1)$, into a square matrix of dimension $n + 1$, the problem is then reduced to the linear case. A fast algorithm based on (Hernández-Toledo 2005) has been implemented in Jarrah et al. (2007), using the symbolic computation package *Macaulay2*.

It is not surprising that the phase space structure of Φ should depend on invariants of the matrix $A = (a_{ij})$. The rational canonical form of A is a block-diagonal matrix, and one can recover the structure of the phase space of A from that of the blocks in the rational form of A . Each block

represents either an invertible or a nilpotent linear transformation. Consider an invertible block B . If $\mu(x)$ is the minimal polynomial of B , then there exists s such that $\mu(x)$ divides $x^s - 1$. Hence $B^s - I = 0$ which implies that $B^s \mathbf{v} = \mathbf{v}$. That is, every state vector $\mathbf{v}$ in the phase space of B is in a cycle whose length is a divisor of s .

Definition 7 For any polynomial $\lambda(x)$ in $k[x]$, the *order* of $\lambda(x)$ is the least integer s such that $\lambda(x)$ divides $x^s - 1$.

The cycle structure of the phase space of Φ can be completely determined from the orders of the irreducible factors of the minimal polynomial of Φ . The computation of these orders involves in particular the factorization of numbers of the form $q^r - 1$, which makes the computation of the order of a polynomial potentially quite costly. The nilpotent blocks in the decomposition of A determine the tree structure at the nodes of the limit cycles. It turns out that all trees at all periodic nodes are identical. This generalizes a result in Martin et al. (1984) for additive cellular automata over the field with two elements.

While the fact that the structure of the phase space of a linear system can be determined from the invariants associated with its matrix may not be unexpected, it is a beautiful example of how the right mathematical viewpoint provides powerful tools to completely solve the problem of relating the structure of the local functions to the resulting (or emerging) dynamics. Linear and affine systems have been studied extensively in several different contexts and from several different points of view, in particular the case of cellular automata. For instance, additive cellular automata over more general rings as state sets have been studied, e.g., in Chaudhuri (1997). Further results on additive CAs can also be found there. One important focus in Chaudhuri (1997) is on the problem of finding CAs with limit cycles of maximal length for the purpose of constructing pseudo random number generators.

Unfortunately, the situation is more complicated for nonlinear systems. For the special class of Boolean synchronous systems whose local update functions consist of monomials, there is a

polynomial time algorithm that determines whether a given monomial system has only fixed points as periodic points (Colón-Reyes et al. 2004). This question was motivated by applications to the modeling of biochemical networks. The criterion is given in terms of invariants of the dependency graph Y . For a strongly connected directed graph Y (i.e., there is a directed path between any pairs of vertices), its *loop number* is the greatest common divisor of all directed loops at a particular vertex. (This number is independent of the vertex chosen.)

Theorem 8 (Colón-Reyes et al. 2004) A Boolean monomial system has only fixed points as periodic points if and only if the loop number of every strongly connected component of its dependency graph is equal to 1.

In Colón-Reyes et al. (2006), it is shown that the problem for general finite fields can be reduced to that of a Boolean system and a linear system over rings of the form $\mathbf{Z}/p'\mathbf{Z}$, p prime. Boolean monomial systems have been studied before in the cellular automaton context (Bartlett and Garzon 1993).

Sequential Update Systems

The update order in a sequential dynamical system has been studied using combinatorial and algebraic techniques. A natural question to ask here is how the system depends on the update schedule. In Barrett et al. (2000, 2001b), Mortveit and Reidys (2001), Reidys (1998), this was answered on several levels for the special case where the update schedule is a permutation. We describe these results in some detail. Results about the more general case of update orders described by words on the indices of the local update functions can be found in Garcia et al. (2006).

Given local update functions $f_i: k^n \rightarrow k$ and permutation update orders σ, π , a natural question is when the two resulting SDS Φ_σ and Φ_π are identical and, more generally, how many different systems one obtains by varying the update order over all permutations. Both questions can be answered in great generality. The answer involves invariants of two graphs, namely, the acyclic orientations of

the dependency graph Y of the local update functions and the *update graph* of Y . The update graph $U(Y)$ of Y is the graph whose vertex set consists of all permutations of the vertex set of Y (Reidys 1998). There is an (undirected) edge between two permutations $\sigma = (\sigma_1, \dots, \sigma_n)$ and $\tau = (\tau_1, \dots, \tau_n)$ if they differ by a transposition of two adjacent entries σ_i and σ_{i+1} such that there is no edge in Y between σ_i and σ_{i+1} .

The update graph encodes the fact that one can commute two local update functions f_i and f_j without affecting the end result Φ if i and j are not connected by an edge in Y . That is, $\dots f_i \circ f_j \dots = \dots f_j \circ f_i \dots$ if and only if i and j are not connected by an edge in Y .

All permutations belonging to the same connected component in $U(Y)$ give identical SDS maps. The number of (connected) components in $U(Y)$ is therefore an upper bound for the number of functionally inequivalent SDS that can be generated by just changing the update order. It is convenient to introduce an equivalence relation $\sim_Y$ on S_Y by $\pi \sim_Y \sigma$ if π and σ belong to the same connected component in the graph $U(Y)$. It is then clear that if $\pi \sim_Y \sigma$, then corresponding sequential dynamical systems are identical as maps. This can also be characterized in terms of acyclic orientations of the graph Y : Each component in the update graph induces a unique acyclic orientation of the graph Y . Moreover, we have the following result:

Proposition 9 (Reidys 1998) There is a bijection

$$F_Y : S_Y / \sim_Y \rightarrow \text{Acyc}(Y),$$

where $S_Y / \sim_Y$ denotes the set of equivalence classes of $\sim_Y$ and $\text{Acyc}(Y)$ denotes the set of acyclic orientations of Y .

This upper bound on the number of functionally different systems has been shown in Reidys (1998) to be sharp for Boolean systems, in the sense that for a given Y one constructs this number of different systems, using appropriate combinations of NOR functions.

For two permutations σ and τ , it is easy to determine if they give identical SDS maps: One

can just compare their induced acyclic orientations. The number of acyclic orientations of the graph Y tells how many functionally different SDS maps one can obtain for a fixed graph and fixed vertex functions. The work of Cartier and Foata (1969) on partially commutative monoids studies a similar question, but their work is not concerned with finite dynamical systems.

Note that permutation update orders have been studied sporadically in the context of cellular automata on circle graphs (Park et al. 1986) but not in a systematic way, typically using the order $(1, 2, \dots, n)$ or the even-odd/odd-even orders. As a side note, we remark that this work also confirms our findings that switching from a parallel update order to a sequential order turns the complex behavior found in Wolfram's "class III and IV" automata into much more regular or mundane dynamics (see, e.g., Schönfisch and de Roos 1999).

The work on functional equivalence was extended to dynamical equivalence (topological conjugation) in Barrett et al. (2001b), Mortveit and Reidys (2001). The automorphism group of the graph Y can be made to act on the components of the update graph $U(Y)$ and therefore also on the acyclic orientations of Y . All permutations contained in components of an orbit under $\text{Aut}(Y)$ give rise to dynamically equivalent sequential dynamical systems, that is, to isomorphic phase spaces. However, here one needs some more technical assumptions, i.e., the local functions must be symmetric and induced (see Barrett et al. 2003c). This of course also leads to a bound for the number of dynamically inequivalent systems that can be generated by varying the update order alone. Again, this was first done for permutation update orders. The theory was extended to words over the vertex set of Y in Garcia et al. (2006), Reidys (2006).

The structure of the graph Y influences the dynamics of the system. As an example, graph invariants such as the independent sets of Y turn out to be in a bijective correspondence with the invariant set of sequential systems over the Boolean field $k = \{0, 1\}$ that have $\text{nor}_t: k^t \rightarrow k$ given by $\text{nor}_t: (x_1, \dots, x_t) = (1 + x_1) \cdots (1 + x_t)$ as local functions (Reidys 2001). This can be extended to other classes such as those with order independent

invariant sets as in Hansson et al. (2005). We have already seen how the automorphisms of a graph give rise to equivalence (Mortveit and Reidys 2001). Also, if the graph Y has nontrivial covering maps, we can derive simplified or reduced (in an appropriate sense) versions of the original SDS over the image graphs of Y (see, e.g., Mortveit and Reidys 2004; Reidys 2005).

Parallel and sequential dynamical systems differ when it comes to invertibility. Whereas it is generally computationally intractable to determine if a CA over $\mathbb{Z}^d$ is invertible for $d \geq 2$ (Kari 2005), it is straightforward to determine this for a sequential dynamical system (Mortveit and Reidys 2001). For example, it turns out that the only invertible Boolean sequential dynamical systems with symmetric local functions are the ones where the local functions are either the parity function or the logical complement of the parity function (Barrett et al. 2001b).

Some classes of sequential dynamical systems such as the ones induced by the nor -function have desirable stability properties (Hansson et al. 2005). These systems have minimal invariant sets (i.e., periodic states) that do not depend on the update order. Additionally, these invariant sets are stable with respect to configuration perturbations.

If a state $\mathbf{c}$ is perturbed to a state $\mathbf{c}'$ that is not periodic, this state will evolve to a periodic state $\mathbf{c}''$ in one step; that is, the system will quickly return to the invariant set. However, the states $\mathbf{c}$ and $\mathbf{c}''$ may not necessarily be in the same periodic orbit.

The Category of Sequential Dynamical Systems

As a general principle, in order to study a given class of mathematical objects, it is useful to study transformations between them. In order to provide a good basis for a mathematical analysis, the objects and transformations together should form a *category*, that is, the class of transformations between two objects should satisfy certain reasonable properties (see, e.g., Mac Lane 1998). Several proposed definitions of a transformation of SDS have been published, notably in Laubenbacher and Pareigis (2003) and Reidys (2005). One possible interpretation of a

transformation of SDS from the point of view of agent-based simulation is that the transformation represents the approximation of one simulation by another or the embedding/projection of one simulation into/onto another. These concepts have obvious utility when considering different simulations of the same complex system.

One can take different points of view in defining a transformation of SDS. One approach is to require that a transformation is compatible with the defining structural elements of an SDS, that is, with the dependency graph, the local update functions, and the update schedule. If this is done properly, then one should expect to be able to prove that the resulting transformation induces a transformation at the level of phase spaces. That is, transformations between SDS should preserve the local and global dynamic behavior. This implies that transformations between SDS lead to transformations between the associated global update functions.

Since the point of view of SDS is that global dynamics emerges from system properties that are defined locally, the notion of SDS transformation should focus on the local structure. This is the point of view taken in Laubenbacher and Pareigis (2003). The definition given there is rather technical and the details are beyond the scope of this entry. The basic idea is as follows. Let $\Phi_\pi = f_{\pi(n)} \circ \dots \circ f_{\pi(1)}$ and $\Phi_\sigma = g_{\sigma(m)} \circ \dots \circ g_{\sigma(1)}$ with the dependency graphs Y_π and Y_σ , respectively. A transformation $F: \Phi_\pi \rightarrow \Phi_\sigma$ is determined by the following:

- A graph mapping $\varphi: Y_\pi \rightarrow Y_\sigma$ (reverse direction)
- A family of maps $k_{\varphi(v)}: k_v \rightarrow k_v$ with $v \in Y_\pi$
- An order preserving map $\sigma \rightarrow \pi$ of update schedules

These maps are required to satisfy the property that they “locally” assemble to a coherent transformation. Using this definition of transformation, it is shown (Theorem 2.6 in Laubenbacher and Pareigis 2003) that the class of SDS forms a category. One of the requirements, for instance, is that the composition of two transformations is again a transformation. Furthermore, it is shown (Theorem 3.2 in

Laubenbacher and Pareigis 2003) that a transformation of SDS induces a map of directed graphs on the phase spaces of the two systems. That is, a transformation of the local structural elements of SDS induces a transformation of global dynamics. One of the results proven in Laubenbacher and Pareigis (2003) is that every SDS can be decomposed uniquely into a direct product (in the categorical sense) of indecomposable SDS.

Another possible point of view is that a transformation

$$F: (\Phi_\pi: k^n \rightarrow k^n) \rightarrow (\Phi_\sigma: k^m \rightarrow k^m)$$

is a function $F: k^n \rightarrow k^m$ such that $F \circ \Phi_\pi = \Phi_\sigma \circ F$, without requiring specific structural properties. This is the approach in Reidys (2005). This definition also results in a category and a collection of mathematical results. Whatever definition chosen, much work remains to be done in studying these categories and their properties.

Future Directions

Agent-based computer simulation is an important method for modeling many complex systems, whose global dynamics emerges from the interaction of many local entities. Sometimes this is the only feasible approach, especially when available information is not enough to construct global dynamic models. The size of many realistic systems, however, leads to computer models that are themselves highly complex, even if they are constructed from simple software entities. As a result, it becomes challenging to carry out verification, validation, and analysis of the models, since these consist in essence of complex computer programs. This entry argues that the appropriate approach is to provide a formal mathematical foundation by introducing a class of mathematical objects to which one can map agent-based simulations. These objects should capture the key features of an agent-based simulation and should be mathematically rich enough to allow the derivation of general results and techniques. The mathematical setting of dynamical systems is a natural choice for this purpose.

The class of finite dynamical systems over a state set X which carries the structure of a finite field satisfies all these criteria. Parallel, sequential, and stochastic versions of these are rich enough to serve as the mathematical basis for models of a broad range of complex systems. While finite dynamical systems have been studied extensively from an experimental point of view, their mathematical theory should be considered to be in its infancy, providing a fruitful area of research at the interface of mathematics, computer science, and complex systems theory.

Bibliography

Primary Literature

- Bagrodia RL (1998) Parallel languages for discrete-event simulation models. *IEEE Comput Sci Eng* 5(2):27–38
- Barrett CL, Reidys CM (1999) Elements of a theory of simulation I: sequential CA over random graphs. *Appl Math Comput* 98:241–259
- Barrett CL, Mortveit HS, Reidys CM (2000) Elements of a theory of simulation II: sequential dynamical systems. *Appl Math Comput* 107(2–3):121–136
- Barrett CL, Hunt III HB, Marathe MV, Ravi SS, Rosenkrantz DJ, Stearns RE, Tosic P (2001) Garden of eden and fixed point configurations in sequential dynamical systems. In: *Proceedings of international conference on combinatorics, computation and geometry DM-CCG*. Paris, France, pp 95–110
- Barrett CL, Mortveit HS, Reidys CM (2001b) Elements of a theory of simulation III: equivalence of SDS. *Appl Math Comput* 122:325–340
- Barrett CL, Marathe MV, Smith JP, Ravi SS (2002) A mobility and traffic generation framework for modeling and simulating ad hoc communication networks. In: *SAC'02*. ACM, Madrid, pp 122–126
- Barrett CL, Hunt HB III, Marathe MV, Ravi SS, Rosenkrantz DJ, Stearns RE (2003a) On some special classes of sequential dynamical systems. *Ann Comb* 7(4):381–408
- Barrett CL, Hunt HB III, Marathe MV, Ravi SS, Rosenkrantz DJ, Stearns RE (2003b) Reachability problems for sequential dynamical systems with threshold functions. *Theor Comput Sci* 295(1–3):41–64
- Barrett CL, Mortveit HS, Reidys CM (2003c) Elements of a theory of computer simulation. IV. Sequential dynamical systems: fixed points, invertibility and equivalence. *Appl Math Comput* 134(1):153–171
- Barrett CL, Hunt HB III, Marathe MV, Ravi SS, Rosenkrantz DJ, Stearns RE (2006) Complexity of reachability problems for finite discrete sequential dynamical systems. *J Comput Syst Sci* 72:1317–1345
- Barrett CL, Hunt III HB, Marathe MV, Ravi SS, Rosenkrantz DJ, Stearns RE, Thakur M (2007) Computational aspects of analyzing social network dynamics. In: *Proceedings of international joint conference on artificial intelligence IJCAI*. Paris, France, pp 2268–2273
- Barrett CL, Hunt HB III, Marathe MV, Ravi SS, Rosenkrantz DJ, Stearns RE, Thakur M (2007b) Predecessor existence problems for finite discrete dynamical systems. *Theor Comput Sci* 1–2:3–37
- Bartlett R, Garzon M (1993) Monomial cellular automata. *Complex Syst* 7(5):367–388
- Bernaschi M, Castiglione F (2002) Selection of escape mutants from immune recognition during hiv infection. *Immunol Cell Biol* 80:307–313
- Bernaschi M, Succi S, Castiglione F (2000) Large-scale cellular automata simulations of the immune system response. *Phys Rev E* 61:1851–1854
- Booch G, Rumbaugh J, Jacobson I (2005) Unified modeling language user guide, 2nd edn. Addison-Wesley, Reading
- Brand D, Zafropulo P (1983) On communicating finite-state machines. *J ACM* 30:323–342
- Cartier P, Foata D (1969) *Problèmes combinatoires de commutation et réarrangements*, vol 85, *Lecture Notes in Mathematics*. Springer, Berlin
- Castiglione F, Agur Z (2003) Analyzing hypersensitivity to chemotherapy in a cellular automata model of the immune system. In: Preziosi L (ed) *Cancer modeling and simulation*. Chapman and Hall/CRC, London
- Castiglione F, Bernaschi M, Succi S (1997) Simulating the immune response on a distributed parallel computer. *Int J Mod Phys C* 8:527–545. <https://doi.org/10.1142/S0129183197000424>
- Castiglione F, Duca K, Jarrah A, Laubenbacher R, Hochberg D, Thorley-Lawson D (2007) Simulating Epstein-Barr virus infection with C-ImmSim. *Bioinformatics* 23(11):1371–1377
- Celada F, Seiden P (1992a) A computer model of cellular interactions in the immune system. *Immunol Today* 13(2):56–62
- Celada F, Seiden P (1992b) A model for simulating cognate recognition and response in the immune system. *J Theor Biol* 158:235–270
- Celada F, Seiden P (1996) Affinity maturation and hypermutation in a simulation of the humoral immune response. *Eur J Immunol* 26(6):1350–1358
- Chaudhuri PP (1997) *Additive cellular automata. Theory and applications*, vol 1. IEEE Computer Society Press, Washington, DC
- Colón-Reyes O, Laubenbacher R, Pareigis B (2004) Boolean monomial dynamical systems. *Ann Comb* 8:425–439
- Colón-Reyes O, Jarrah A, Laubenbacher R, Sturfels B (2006) Monomial dynamical systems over finite fields. *Complex Syst* 16(4):333–342
- Dawson D (1974) Synchronous and asynchronous reversible Markov systems. *Canad Math Bull* 17(5):633–649
- Ebeling W, Schweitzer F (2001) Swarms of particle agents with harmonic interactions. *Theor Biosci* 120–3(4):207–224
- Elsaspas B (1959) The theory of autonomous linear sequential networks. *IRE Trans Circuit Theor* 1:45–60

- Farmer J, Packard N, Perelson A (1986) The immune system, adaptation, and machine learning. *Phys D* 2(1–3):187–204
- Frish U, Hasslacher B, Pomeau Y (1986) Lattice-gas automata for the Navier-Stokes equations. *Phys Rev Lett* 56:1505–1508
- Fukš H (2004) Probabilistic cellular automata with conserved quantities. *Nonlinearity* 17:159–173
- Garcia LD, Jarrah AS, Laubenbacher R (2006) Sequential dynamical systems over words. *Appl Math Comput* 174(1):500–510
- Gardner M (1970) The fantastic combinations of John Conway's new solitaire game "life". *Sci Am* 223:120–123
- Gouda M, Chang C (1986) Proving liveness for networks of communicating finite-state machines. *ACM Trans Program Lang Syst* 8:154–182
- Guo Y, Gong W, Towsley D (2000) Time-stepped hybrid simulation (TSHS) for large scale networks. In: *INFOCOM 2000. Proceedings of nineteenth annual joint conference of the IEEE computer and communications societies*, vol 2. IEEE, Washington, DC, pp 441–450
- Gupta A, Katiyar V (2005) Analyses of shock waves and jams in traffic flow. *J Phys A* 38:4069–4083
- Hansson ÅÅ, Mortveit HS, Reidys CM (2005) On asynchronous cellular automata. *Adv Complex Syst* 8(4):521–538
- Hedlund G (1969) Endomorphisms and automorphisms of the shift dynamical system. *Math Syst Theory* 3:320–375
- Hernández-Toledo A (2005) Linear finite dynamical systems. *Commun Algebra* 33(9):2977–2989
- Hopcroft JE, Motwani R, Ullman JD (2000) *Automata theory, languages and computation*. Addison Wesley, Reading
- Hopfield J (1982) Neural networks and physical systems with emergent collective computational properties. *Proc Natl Acad Sci U S A* 79:2554–2588
- Ilachinsky A (2001) *Cellular automata: a discrete universe*. World Scientific, Singapore
- Jarrah A, Laubenbacher R, Stillman M, Vera-Licona P (2007) An efficient algorithm for the phase space structure of linear dynamical systems over finite fields (submitted)
- Jefferson DR (1985) Virtual time. *ACM Trans Program Lang Syst* 7(3):404–425
- Kari J (2005) Theory of cellular automata: a survey. *Theory Comput Sci* 334:3–33
- Keyfitz BL (2004) Hold that light! Modeling of traffic flow by differential equations. *Stud Math Libr* 26:127–153
- Kozen DC (1997) *Automata and computability*. Springer, New York
- Laubenbacher R, Pareigis B (2003) Decomposition and simulation of sequential dynamical systems. *Adv Appl Math* 30:655–678
- Lidl R, Niederreiter H (1997) *Finite fields*. Cambridge University Press, Cambridge
- Liggett TM (2005) *Interacting particle systems*. Classics in mathematics. Springer, Berlin, Reprint of the 1985 original
- Lind DA (1984) Applications of ergodic theory and sofic systems to cellular automata. *Phys D* 10D:36–44
- Lindgren K, Moore C, Nordahl M (1998) Complexity of two-dimensional patterns. *J Stat Phys* 91(5–6):909–951
- Mac Lane S (1998) *Category theory for the working mathematician*, 2nd edn. Springer, New York, No 5. in GTM
- Macy MW, Kitts JA, Flache A (2003) Polarization in dynamic networks: a Hopfield model of emergent structure. In: *Dynamic social network modeling and analysis*. The National Academies Press, Washington, DC, pp 162–173
- Martin O, Odlyzko A, Wolfram S (1984) Algebraic properties of cellular automata. *Commun Math Phys* 93:219–258
- Milligan D, Wilson M (1993) The behavior of affine Boolean sequential networks. *Connect Sci* 5(2):153–167
- Minar N, Burkhart R, Langton C, Manor A (1996) The swarm simulation system: a toolkit for building multi-agent simulations. Santa Fe Institute preprint series. <http://www.santafe.edu/research/publications/wpabstract/199606042>. Accessed 11 Aug 2008
- Misra J (1986) Distributed discrete-event simulation. *ACM Comput Surv* 18(1):39–65
- Moncion T, Hutzler G, Amar P (2006) Verification of biochemical agent-based models using petri nets. In: Robert T (ed) *International symposium on agent based modeling and simulation, ABModSim'06*. Austrian Society for Cybernetics Studies, pp 695–700. <http://www.ibisc.univ-evry.fr/pub/basilic/OUT/Publications/2006/MHA06>
- Morpurgo D, Serenita R, Seiden P, Celada F (1995) Modelling thymic functions in a cellular automaton. *Int Immunol* 7:505–516
- Mortveit HS, Reidys CM (2001) Discrete, sequential dynamical systems. *Discret Math* 226:281–295
- Mortveit HS, Reidys CM (2004) Reduction of discrete dynamical systems over graphs. *Adv Complex Syst* 7(1):1–20
- Nagel K, Schreckenberg M (1992) A cellular automaton model for freeway traffic. *J Phys I* 2:2221–2229
- Nagel K, Wagner P (2006) *Traffic flow: approaches to modelling and control*. Wiley, Hoboken, NJ
- Nagel K, Schreckenberg M, Schadschneider A, Ito N (1995) Discrete stochastic models for traffic flow. *Phys Rev E* 51:2939–2949
- Nagel K, Rickert M, Barrett CL (1997) Large-scale traffic simulation, vol 1215, *Lecture notes in computer science*. Springer, Berlin, pp 380–402
- Nance RE (1993) A history of discrete event simulation programming languages. *ACM SIGPLAN Not* 28:149–175
- North MJ, Collier NT, Vos JR (2006) Experiences creating three implementations of the repast agent modeling toolkit. *ACM Trans Model Comput Simul* 16:1–25
- Oroponen P (1994) Computational complexity of neural networks: a survey. *Nord J Comput* 1:94–110
- Oroponen P (1996) The computational power of discrete hopfield networks with hidden units. *Neural Comput* 8:403–415

- Park JK, Steiglitz K, Thruston WP (1986) Soliton-like behavior in automata. *Phys D* 19D:423–432
- Reidys C (1998) Acyclic orientations of random graphs. *Adv Appl Math* 21:181–192
- Reidys CM (2001) On acyclic orientations and sequential dynamical systems. *Adv Appl Math* 27:790–804
- Reidys CM (2005) On certain morphisms of sequential dynamical systems. *Discret Math* 296(2–3):245–257
- Reidys CM (2006) Sequential dynamical systems over words. *Ann Comb* 10(4):481–498
- Rickert M, Nagel K, Schreckenberg M, Latour A (1996) Two lane traffic simulations using cellular automata. *Phys A* 231:534–550
- Rothman DH (1988) Cellular-automaton fluids: a model for flow in porous media. *Geophysics* 53:509–518
- Russell S, Norwig P (2003) Artificial intelligence: a modern approach. Prentice-Hall, Upper Saddle River
- Schönfisch B, de Roos A (1999) Synchronous and asynchronous updating in cellular automata. *BioSystems* 51:123–143
- Shmulevich I, Dougherty ER, Kim S, Zhang W (2002a) Probabilistic boolean networks: a rule-based uncertainty model for gene regulatory networks. *Bioinformatics* 18(2):261–274
- Shmulevich I, Dougherty ER, Zhang W (2002b) From boolean to probabilistic boolean networks as models of genetic regulatory networks. *Proc IEEE* 90(11):1778–1792
- Sipser M (1997) Introduction to the theory of computation. PWS Publishing Co, Boston
- Vasershtein L (1969) Markov processes over denumerable products of spaces describing large system of automata. *Probl Peredachi Inf* 5(3):64–72
- von Neumann J, Burks AW (eds) (1966) Theory of self-reproducing automata. University of Illinois Press, Champaign
- Whitham G (1999) Linear and nonlinear waves, reprint edition edn. Pure and applied mathematics: a Wiley-Interscience series of texts, monographs and tracts. Wiley-Interscience, New York
- Wolfram S (1983) Statistical mechanics of cellular automata. *Rev Mod Phys* 55:601–644
- Wolfram S (1986) Theory and applications of cellular automata, vol 1, Advanced series on complex systems. World Scientific Publishing Company, Singapore
- Wolfram S (2002) A new kind of science. Wolfram Media, Champaign Books and Reviews, Champaign
- Wooldridge M (2002) Introduction to multiagent systems. Wiley, Chichester



Logic and Geometry of Agents in Agent-Based Modeling

Samson Abramsky

Department of Computer Science, University of Oxford, Oxford, UK

Article Outline

Definition of the Subject

Introduction

Toward a Logic and Geometry of Interaction

Discussion

Further Directions

Bibliography

Glossary

Agent A computational or biological entity which can perform actions which affect its environment and observe actions performed by the environment or other agents.

Agent-based modeling Modeling a complex system in terms of multiple agents interacting with each other.

Combinators Basic operations in a function algebra. Combinatory logic is a variable-free formulation of functional computation.

Compositionality Describing a complex system or object in a structured fashion, as built up by applying certain operations hierarchically, starting from a stock of basic types of system. Compositional definitions of functions of systems are those which respect this hierarchical structure.

Information flow The process whereby information held by one agent or part of a system is transferred to another, possibly in some modified form. Information flow, which is generally mediated by interaction, is caused by agents performing actions and other agents observing those actions.

Interaction A pattern of actions in a multiagent system. Each action is performed by some agent and may be observed by others.

Linear logic A substructural logic in which the operations of copying and deleting premises are not allowed in general.

Definition of the Subject

Agent-based modeling has become of increasing importance in computer science and also in mathematical modeling and simulation. The idea is that the behavior of a complex system can be described as arising from the interaction of multiple agents, with each other and with the environment, using simple local rules. It is widely recognized that building a sound and widely applicable theory for such systems will require an interdisciplinary approach and the development of new mathematical and computational concepts.

In this article, agents and interaction will be studied from the perspective of logic and computer science. It will be shown how ideas about logical dynamics, games, and geometry of interaction, which have been developed over the past two decades, lead toward a structural theory of agents and interaction. This provides a basis for powerful logical methods such as compositionality, types, and high-level calculi, which have proved so fruitful in computer science, to be applied in this domain.

This approach should be contrasted with the more familiar approaches to agent-based modeling using primitive agents such as cellular automata. The main focus of such approaches is systems modeling, where the agents are essentially discrete counterparts to traditional PDE or ODE models of dynamical systems. Ultimately, these approaches should be seen as complementary, and the goal is to combine them. In the present article, the emphasis is on presenting some basic ideas of the novel approach to complex systems using logical and compositional methods.

Introduction

In this article, a logical and geometric perspective on the information flow arising from agent interaction will be described. This is a novel and distinctive approach within the emerging field of agent-based modeling. While it is still in its formative stages, there are already some key structural insights which it offers which will be fundamental to any deep and comprehensive theory of agents.

Firstly, a fundamental methodological point will be discussed, which has played a crucial role in computer science for several decades but has yet to achieve the recognition in general scientific modeling which it deserves: the importance of **compositionality**.

Compositionality

A methodological principle from computer science (and logic) of **major** potential importance for mathematical modeling throughout the sciences:

- **Traditional approach:** Whole-system (monolithic) analysis of given systems. A key role is played by structuring templates, e.g., “Find the Hamiltonian.”
- **Compositional approach:** Start with a fixed set of basic (simple) building blocks and **constructions** for building new (in general more complex) systems out of given subsystems and build up the required complex system with these.

More formally, compositionality can be expressed **algebraically**:

$$S = \omega(S_1, \dots, S_n).$$

The system S is described as being built up from subsystems $S_1, \dots, S_n$ by the operation ω . There is also a **logical** perspective:

$$\frac{S_1 \models \phi_1, \dots, S_n \models \phi_n}{\omega(S_1, \dots, S_n) \models \phi}.$$

(Read $S \models \phi$ as “system S satisfies the property ϕ .”) Here **properties** ϕ of the compound system S can be inferred by verifying properties $\phi_1, \dots, \phi_n$ for the simpler subsystems $S_1, \dots, S_n$.

Some Key Points

$$S = \omega_1(\omega_2(a_1, a_2, a_3); \omega_1(a_4, a_5)).$$

- The compositional view of complex systems sees them as built up, not just by one-level composition of basic agents, $S = \parallel_{i \in I} a_i$, which is the usual scenario in current agent-based modeling and simulation, but **hierarchically**:

It tracks the properties of the sub-compound systems all the way up (or down) the tree of syntax.

- The **available repertoire** of system constructors is also an important aspect of the modeling here, leading to questions of **expressiveness** and **functional completeness**.

This paradigm has played a major role in computer science over the past four decades and is already starting to be applied in quantum computing (quantum programming languages), biological modeling (process calculi), and business modeling (idem) and will surely be applied more widely and deeply in economics as well as physical and biological sciences.

This kind of modeling carries in its train a range of powerful **analytical techniques**: types, semantics, verification, model checking, etc.

Computation as Interaction

The second key point to be emphasized is the emergence of **interaction** as a key notion in computation and, increasingly in a wide range of scientific fields, in physics (quantum information), biology (at a range of scales and levels, from molecular interactions to evolutionary theory), economics (game theory), linguistics (dialogical analysis), etc. Moreover, these developments are feeding back into logic and philosophy.


Note that interaction happens **between agents** – a theory of interaction amounts to an **agent dynamics**.

Changing Views of Computation

The scene can be set by recalling how perspectives on computation have changed since the first computers appeared. The early practice of

Logic and Geometry of Agents in Agent-Based Modeling,

Fig. 1 Computing “in the isolation ward”



computing is shown in Fig. 1. This is the era of stand-alone machines and programs: computers are served by an elite priesthood and have only a narrow input–output interface with the rest of the world.


First-Generation Models of Computation These models live on the existing intellectual inheritance from discrete mathematics and logic. **Time** and **processes** lurk in the background but are largely suppressed.

Given this limited vision of computing, there is a very natural abstraction of computation, in which programs are seen as computing **functions** or **relations** from inputs to outputs.



Computation in the age of the Internet


Instead of isolated systems, with rudimentary interactions with their environment, the standard unit of description or design becomes a **process** or **agent**, the essence of whose behavior is **how it interacts** with its environment.



The interaction between the system and the environment can be conceptualized as a **two-person game**. A program specifying how the system should behave in the face of all possible actions by the environment is then a **strategy** for the player corresponding to the system.

Interaction Complex behavior arises as the global effect of a **system** of **interacting agents** (or processes).

The key building block is the agent. The key operation is **interaction** – plugging agents together so that they interact with each other.



Who is the system? Who is the environment? This **symmetry** between the system and the environment carries a first clue that there is some structure here; it will lead to a key **duality** and a deep connection to logic.

Toward a Compositional Approach to Complex Systems This conceptual model works at all “scales”:

- Macro-scale: processes in operating systems, software agents on the Internet, and transactions
- Microscale: how programs are implemented (subroutine call-return protocols, register transfer) all the way down into the hardware

It is applicable both to **design** (synthesis) and to **description** (analysis) and to **artificial** and to **natural** information-processing systems.

There are of course large issues lurking here, e.g., in the realm of “complex systems”: **emergent behavior** and even **intelligence**. Is it helpful, or even feasible, to understand this complexity **compositionally**? New conceptual tools and new theories are needed to help us analyze and synthesize these systems, **understand**, and **build**.

Toward a Logic and Geometry of Interaction

Toward a “Logic of Interaction”

Specifying and reasoning about the behavior of computer programs take us into the realm of logic. For the first-generation models, logic could be taken “as it was” – static and timeless. For the second-generation models, getting an adequate account – a genuine “logic of interaction” – may require a fundamental reconceptualization of logic itself. This radical revision of the view of logic is happening anyway – prompted partly by the applications and partly by ideas arising within logic.

The Static Conception of Logic

The usual “static” notion of tautology is as “a statement which is vacuously true because it is compatible with all states of affairs”:

$$A \vee \neg A.$$

“It is raining **or** it is not raining” – truth-functional semantics. This is illustrated (subversively) in Fig. 2. But what could a **dynamic notion of tautology** look like?

The Copycat Strategy

Consider the following little fable, illustrated in Fig. 3:

How to beat an International Chess Grandmaster by the power of pure logic.

The idea is to rely on logic, rather than on any talent at chess. We arrange to play two games of



Logic and Geometry of Agents in Agent-Based Modeling, Fig. 2 Tertium non datur?



Logic and Geometry of Agents in Agent-Based Modeling, Fig. 3 How to beat a Grandmaster

chess with the Grandmaster, say Gary Kasparov, once as White and once as Black. Moreover, we so arrange matters that we start with the game in which we play as Black. Kasparov makes his opening move; we respond by playing the **same** move in the **other** game – this makes sense, since we are playing as White there. Now Kasparov responds (as Black) to our move in that game; and we copy that response back in the first game. We simply proceed in this fashion, copying the moves that our opponent makes in one board to the other board. The net effect is that **we play the same game twice – once as White and once as Black**. (We have essentially made Kasparov play against himself.) Thus, whoever wins that game, we can claim a win in one of our games against Kasparov! (Even if the game results in a stalemate, we have done as well as Kasparov over the two games – surely still a good result!) Of course, this idea has nothing particularly to do with chess. It can be applied to any two-person game of a very general form. The use of chessboards to illustrate the discussion will continue, but this underlying generality should be kept in mind.

What are the salient features which can be extracted from this example?

A dynamic tautology

There is a sense (which will shortly be made more precise) in which the copycat strategy can be seen as a **dynamic version** of the tautology $A \rightarrow \vee A$. Note, indeed, that an essential condition for being able to play the copycat is that the roles of the two players are interchanged on one board as compared to the other. Note also the disjunctive quality of the argument that we must win in one or other of the two games. But the copycat strategy is a **dynamic process**: a two-way channel which maintains the correlation between the plays in the two games.

Conservation of information flow

The copycat strategy does not **create** any information; it reacts to the environment in such a way that information is conserved. It ensures that exactly the same information flows out to the environment as what flows in. Thus one gets a sense of logic appearing in the form of **conservation laws for information dynamics**.

The power of copying


Another theme which appears here, and of which more will be seen later, concerns the surprising power of simple processes of copying information

from one place to another. Indeed, as shall eventually be seen, such processes are **computationally universal**.

The geometry of information flow

From a dynamical point of view, the copycat strategy realizes a channel between the two game boards, by performing the actions of copying moves. But there is also some implicit geometry here. Indeed, the very idea of two boards laid out side by side appeals to some basic underlying spatial structure. In these terms, the copycat channel can also be understood geometrically, as creating a graphical link between these two spatial locations. These two points of view are complementary and link the logical perspective to powerful ideas arising in modern geometry and mathematical physics.


Further evidence that the copycat strategy embodies more substantial ideas than might at first be apparent can be obtained by varying the scenario. Consider now the case where we play against Kasparov on **three boards**; one as Black and two as White.



Does the copycat strategy still work here? In fact, it can easily be seen that it does **not**. Suppose Kasparov makes an opening move m_1 in the left-hand board where he plays as White; we copy it to the board where we play as White; he responds with m_2 ; and we copy m_2 back to the board where Kasparov opened. So far, all has proceeded as in our original scenario. But now Kasparov has the option of playing a **different** opening move, m_3 say, in the rightmost board. We have no idea how to respond to this move; nor can we copy it anywhere, since the board where we play as White is already “in use.” This shows that these simple

ideas already lead us naturally to the setting of a **resource-sensitive** logic, in which in particular the contraction rule, which can be expressed as $A \rightarrow A \wedge A$ (or equivalently as $\neg A \vee (A \wedge A)$), cannot be assumed to be valid.

What about the other obvious variation, where we play on two boards as White and one as Black?



It seems that the copycat strategy **does** still work here, since we can simply ignore one of the boards where we play as White. However, a geometrical property of the original copycat strategy has been lost, namely, a **connectedness** property that information flows to every part of the system. This at least calls the corresponding logical principle of weakening, which can be expressed as $A \wedge A \rightarrow A$, (or equivalently as $\neg A \vee \neg A \vee A$) into question.

These remarks indicate that we are close to the realm of linear logic and its variants and, mathematically, to the world of monoidal (rather than Cartesian) categories.

Game Semantics

These ideas find formal expression in **game semantics**. Games play the role of:

- Interface types for computation modules
- Propositions with dynamic content

In particular, two-person games capture the duality of:

- Player versus opponent
- System versus environment

Agents are Strategies In this setting, agents or processes can be modeled as **strategies** for playing

the game. These strategies **interact** by playing against each other. A notion of correctness is obtained which is **logical** in character in terms of the idea of **winning** strategy – one which is guaranteed to reach a successful outcome however the environment behaves. This in a sense replaces (or better, **refines**) the logical notion of “truth”: winning strategies are the dynamic version of tautologies (more accurately, of their **proofs**).

Building Complex Systems by Combining Games It will now be seen how games can be combined to produce more complex behaviors while retaining control over the interface. This provides a basis for the **compositional** understanding of systems of interacting agents – understanding the behavior of a complex system in terms of the behavior of its parts. This is crucial for both analysis and synthesis, i.e., for both description and design. These operations for building games can be seen as (dynamic forms of) “type constructors” or “logical connectives.” (The underlying logic here will in fact be linear logic.)

Duality – Linear Negation $A^\perp$ – interchange roles of player and opponent (reflecting the symmetry of interaction).

Note that, with this interpretation, negation is involutive:

$$A^{\perp\perp} = A.$$

Tensor – Linear Conjunction

$$A \otimes B$$

The idea here is to combine the two game boards into one system, **without any information flow between the two subsystems**. (This is the significance of the “wall” separating the two players, who will be referred to as Gary (Kasparov) and Nigel (Short).) This connective has a conjunctive quality, since we must independently be able to play (and to win) in each conjunct. Note, however, that there is no constraint on information flow for the environment, as it plays against this compound system.

Par – Linear Disjunction

$$A \wp B.$$

In this case, there are two boards, but one player (who shall be referred to as the copycat), indicating that we **do** allow information flow for this player between the two game boards. This, for example, allows information revealed in one game board by the opponent to be used against him on the other game board – as exemplified by the copycat strategy. However, note that the wall appears on the environment’s side now, indicating that the environment is constrained to play separately on the two boards, with no communication between them.

Thus there is a De Morgan duality between these two connectives, mediated by the linear negation:

$$(A \otimes B)^\perp = A^\perp \wp B^\perp,$$

$$(A \wp B)^\perp = A^\perp \otimes B^\perp.$$

The idea is that on one side of the mirror of duality (player/system for the tensor, opponent/environment for the par), there is constraint of no information flow, while on the other side, there is information flow.

The copycat strategy can now be reconstrued in logical terms:

It can be seen that it is indeed a winning strategy for $A^\perp \wp A$. Moreover, $A - o B$ (“linear implication”) can be defined by

$$A - o B \equiv A^\perp \wp B,$$

(cf. $A \supset B \equiv \neg A \vee B$.) Then the copycat strategy becomes the canonical proof of the most basic tautology of all: $A - o A$.


The information flow possibilities of par receive a more familiar logical interpretation in terms of the linear implication; namely, that information about the antecedent can be used in proving the consequent (and conversely with respect to their negations, if proof by contraposition is considered).

Thus an entire “linearized” logical structure opens up before us, with a natural interpretation in terms of the dynamics of information flow.

Interaction


A key step in the development now arises: the modeling of **interaction** itself. Constructors create “potentials” for interaction; the operation of plugging modules together so that they can communicate with each other **releases** this potential into **actual computation**. Consider the diagram shown in Fig. 4. Here two separate subsystems are shown, each with a compound structure, expressed by the **logical types of their interfaces**. What these types tell us is that these systems are **composable**; in particular, the **output type** of the first system, namely, B , matches the input type of the second system. Note that this “logical plug compatibility” makes essential use of the duality, just as the copycat strategy did. What makes Gary (the player for the first system), a fit partner for interaction with Nigel (the player for the second system), is that they have **complementary views** of their locus of interaction, namely, B . Gary will play in this type “positively,” as player (he sees it as B), while Nigel will play “negatively,” as opponent (he sees it as $B^\perp$). Thus each will become part of the environment of the other – part of the potential environment of each will be realized by the other – and hence part of the **potential** behavior of each will become **actual** interaction.

This leads to a dynamical interpretation of the fundamental operation of **composition**, in mathematical terms:




or of the **Cut rule**, in logical terms:

$$\text{Cut} : \frac{\vdash \Gamma, A \vdash A^\perp, \Delta}{\Gamma, \Delta}$$



Composition as Interaction

The picture here shows the new system formed by plugging together the two subsystems. The “external interface” to the environment now shows just the left-hand board A as input and the right-hand board C as output. The cut formula B is hidden from the environment and becomes the locus of interaction inside the black box of the system. Suppose that the environment makes



Logic and Geometry of Agents in Agent-Based Modeling, Fig. 4 Two composable systems

some move m in C . This is visible only to Nigel, who as a strategy for B -o C has a response. Suppose this response m_1 is in B . This is a move by Nigel as player in $B^\top$, hence appears to Gary as a move by opponent in B . Gary as a strategy for A -o B has a response m_2 to this move. If this response is again in B , Nigel sees it as a response by the environment to his move and will have a response again and so on. Thus there is a sequence of moves $m_1, \dots, m_k$ in B , ping-ponging back and forth between Nigel and Gary. If, eventually, Nigel responds to Gary's last move by playing in C , or Gary responds to Nigel's last move by playing in A , then this provides the response of the **composed strategy** Gary and Nigel to the original move m . Indeed, all that is visible to the environment is that it played m , and eventually some response appeared in A or C .

Moreover, if both Nigel and Gary are winning strategies, then so is the composed strategy; and the composed strategy will not get stuck forever in the internal ping-pong in B . To see this, suppose for a contradiction that it did in fact get stuck in B . Then there would be an infinite play in B following the winning strategy Gary for **player** in B and the **same** infinite play following the winning strategy Nigel for player in $B^\perp$, hence for **opponent** in B . Hence the same play would count as a win for both player and opponent. This yields the desired contradiction.

Discussion

Game semantics in the sense discussed in this section has had an extensive development over the past decade and a half, with a wealth of applications to the semantics of programming languages, type theories, and logics (Abramsky and Jagadeesan 1994b; Abramsky and McCusker 1997, 1999a; b; Abramsky and Mellies 1999; Abramsky et al. 2000; Hyland and Ong 2000). More recently, there has been an algorithmic turn and some striking applications to verification and program analysis (Abramsky 2002; Abramsky

et al. 2004; Ghica and McCusker 2000; Murawski et al. 2005).

From the point of view of the general analysis of information, there are the following promising lines of development:

- Game semantics provides a promising arena for exploring the combination of quantitative and qualitative theories of information. In particular, it provides a setting for quantifying information flow between agents. Important quantitative questions can be asked about **rate of information flow** through a strategy (representing a program or a proof); how can a system gain **maximum** information from its environment while providing **minimal** information in return; robustness in the presence of **noise**, etc.
- As in the discussion of the copycat strategy, there is an intuition of logical principles arising as **conservation laws for information flow**. (And indeed, in the case of multiplicative linear logic, the proofs correspond exactly to “generalized copycat strategies.”) Can this intuition be developed into a full-fledged theory? Can logical principles be **characterized** as those expressing the conservation principles of this information flow dynamics?
- There is also the hope that the more structured setting of game semantics will usefully constrain the exuberant variety of possibilities offered by process algebra and allow a sharper exploration of the logical space of possibilities for information dynamics. This has already been borne out in part by the success of game semantics in exploring the space of programming language semantics. It has been possible to give crisp characterizations of the “shapes” of computations carried out within certain **programming disciplines**: including purely functional programming (Abramsky et al. 2000; Hyland and Ong 2000), stateful programming (Abramsky and McCusker 1997, 1999a), general references (Abramsky et al. 1998), programming with nonlocal jumps and exceptions (Laird 1997, 2001), non-determinism (Harmer and McCusker 1999),

probability (Danos and Harmer 2002), concurrency (Ghica and Murawski 2004; Ghica and Murawski 2006), names (Abramsky et al. 2004b), polymorphism (Abramsky and Jagadeesan 2005; Hughes 2000), and more. See (Abramsky and McCusker 1999b) for an overview (now rather out of date).

There has also been a parallel line of development of giving **full completeness** results for a range of logics and type theories, characterizing the “space of proofs” for a logic in terms of informatic or geometric constraints which pick out those processes which are proofs for that logic (Abramsky and Jagadeesan 1994b; Abramsky and Mellies 1999; Blute et al. 1998, 2005; Devarajan et al. 1999; Loader 1994). This allows a new look at such issues as the boundaries between classical and constructive logic or the fine structure of polymorphism and second-order quantification.

- This also gives some grounds for optimism that what **computational processes** are can be captured – in a “machine-independent”, and moreover “geometrical”, noninductive way – **without** referring back to Turing machines or any other explicit machine model.
- In the same spirit as for computability, can **polynomial time computation** and other complexity classes be characterized in such terms?

Emergent Logic: The Geometry of Information Flow


The geometric structure inherent in interaction and information flow, and its connections to the dynamic view of interaction, will now be considered. Recent work has shown how common structure arises in a variety of contexts: in **logic** (geometric representations of proofs), **geometry** (diagram algebras, especially the Temperley–Lieb algebra, with connections to knot theory and geometric topology), **computation** (especially functional computation), and **quantum mechanics** (quantum information protocols, exploiting the information flow inherent in quantum entanglement).

A rapid tour will be made through these disparate areas, focusing on how common structure arises. The exposition will rely extensively on

graphical calculi which provide an intuitive and visually appealing window onto the various formalisms to be encountered. These calculi have a substantial mathematical content, founded on the diagrammatic approach to tensor categories; further details can be found in the references.

Logic

Firstly multiplicative linear logic (Girard 1987), the logic of the linear connectives $\otimes$, $\wp$, $(\)^\perp$ which have already been encountered, will be considered as a basic paradigmatic example. A key insight (Girard 1987) is that the essential information in a proof in this system is given by a pairwise matching of the occurrences of positive and negative literals in the sequent – a **proof structure**. For example, the two possible proof structures for the sequent $\alpha^\perp \wp^\perp \alpha^\perp, \alpha \otimes \alpha$ are:



A **geometric** view of proof structures, as certain graphs obtained from the forest of formula trees in the sequent by drawing arcs – “axiom links” – between the paired literal occurrences. Alternatively, a more **dynamic** view can be taken, and proof structures can be represented as **involutions** on the set of literal occurrences in the sequent. These must be **literal preserving**, i.e., an occurrence of a literal l must be mapped to an occurrence of its dual $l^\perp$. Such functions represent a flow of information around the sequent. They can be viewed as **copycat strategies**.

Every sequent proof determines a proof structure. The fundamental question is: **Which proof structures arise from sequent proofs?** A first

answer is **geometric** (or topological) in character. For each proof structure, a **switching graph** can be obtained by deleting, for each occurrence of a subformula $A \wp B$, exactly one of the arcs $A \multimap A \wp B \multimap B$ connecting it to its immediate subformulas. If all such switching graphs are trees, the proof structure is said to be a **proof net**.

A second form of characterization is **interactive**. An orthogonality relation can be defined, between permutations on the set of occurrences of literals in the sequent $\Gamma: f \perp g$ if fg is cyclic. The idea is that f is a candidate proof net, while g is an attempted counterproof – a passage through the literals induced by a choice of switching graph. Note that the alternating composition of f and g expresses **interaction** between f and g , thought of as **strategy** and **counter-strategy**. It generates a path along which information flows around the system.

A semantics of MLL proofs can be given by specifying, for each formula A , a set S of permutations on the set of literal occurrences $|A|$, such that $S = S^{\perp\perp}$, where $S^{\perp} = \{g \mid \forall f \in S: f \perp g\}$. For a literal, the unique permutation (the identity) is specified:

$$S(A \otimes B) = \{f + g \mid f \in S(A) \wedge g \in S(B)\}^{\perp\perp}$$

$$S(A \& B) = S(A^{\perp} \otimes B^{\perp})^{\perp}.$$

Here $f + g$ is a disjoint union of permutations, expressing the absence of information flow (or **information independence**).

Theorem 1 (Sequentializability (Girard 1987) and Full Completeness (Abramsky and Jagadeesan 1994b)) *Let f be a literal-respecting involution on $|\Gamma|$. The following are equivalent: (i) f is the permutation assigned to a sequent proof of Γ ; (ii) f is a proof net; (iii) $f \in S(\Gamma)$.*

This shows that **the geometric and interactive characterizations of the space of proofs coincide**.

Further Developments From a computational perspective, the multiplicative connectives embody **concurrency** and **causal independence**. The scope of the enterprise is greatly expanded

when the other **levels of connective** in linear logic are incorporated:

Additives

The additive conjunction and disjunction allow **causality**, **conflict**, and **case analysis/conditionals** to be expressed. The interaction between the additive and multiplicative levels is rather subtle. The theory sketched above for the multiplicatives has in large part been lifted to multiplicative additive linear logic (Abramsky and Mellies 1999), but a number of key questions and issues needed for a deeper analysis remain to be investigated.

Exponentials

The multiplicative fragment only allows linear time computation to be expressed (under the Curry–Howard paradigm). For a full analysis of computationally expressive systems, it is necessary to allow for **copying** and **deleting**, as regulated by the **exponential connectives** of linear logic. Existing results have extended the multiplicative theory to various systems of typed λ -calculus (corresponding to various forms of intuitionistic multiplicative exponential linear logic) and have begun to investigate systems which, by constraining the exponential types in certain natural ways, capture significant **complexity classes**, especially **P**TIME.

Diagram Algebras

It will now be indicated how this apparently very specialized corner of proof theory in fact connects directly to a broad topic arising in representation theory and knot theory, with connections to mathematical physics. On the one hand, some structure will be **lost**, by obliterating the distinction


between $\otimes$ and $\wp$; this corresponds to moving from ***-autonomous** to **compact closed** categories. This means that the formula tree structure can be dispensed with altogether; it is simply a matter of connecting up literal occurrences, which shall be drawn as “joining up the dots.” Motivation: compact closed categories show up in many contexts of interest!

On the other hand, rather than one-sided sequents, general arrows or two-sided sequents will be represented diagrammatically. This means arrows of the form


$$A_1 \otimes \cdots \otimes A_n \rightarrow B_1 \otimes \cdots \otimes B_m,$$

where each A_i and B_j is a literal, can be represented by involutions on $\{1, \dots, n\} + \{1, \dots, m\}$, which are literal preserving in the extended sense that **opposite** literals are connected in the domain or in the codomain, while occurrences of the **same** literal in the domain and the codomain are connected. An advantage of this representation is that composition is expressed very transparently, by “stacking” arrows.

Example The composition of



is given by




Temperley–Lieb Algebra




The **Temperley–Lieb algebra** played a central role in the **Jones polynomial invariant of knots** (VFR 1985) and ensuing developments. It was originally presented, rather forbiddingly, in terms

of abstract generators and relations (VFR 1985). It was recast in beautifully elementary and conceptual terms by Louis Kauffman as a **planar diagram algebra** (Kauffman 1990).

Generators:



Relations:



The general form of an element of the algebra (actually of the basic multiplicative monoid: the algebra is then constructed freely over this as the “monoid algebra”) is obtained by “joining up the dots” in a **planar** fashion. Multiplication xy is defined by identifying the bottom row of x with the top row of y . In general loops may be formed – these are “scalars,” which can float freely across these figures, represented symbolically by δ above.

How does this connect to knots? A key conceptual insight is due to Kauffman, who saw how to recast the Jones polynomial in elementary combinatorial form in terms of his **bracket polynomial**. The basic idea of the bracket polynomial is expressed by the following equation:

$$\langle \times \rangle = A \langle \smile \rangle + B \langle \frown \rangle \quad \langle \rangle$$

Each overcrossing in a knot or link is evaluated to a weighted sum of the two possible planar smoothings. With suitable choices for the coefficients A and B (as Laurent polynomials), this is invariant under the second and third Reidemeister moves. With an ingenious choice of normalizing factor, it becomes invariant under the first Reidemeister move – and yields the Jones polynomial! What this means algebraically is that the braid group has a representation in the Temperley–Lieb algebra – the above bracket evaluation shows how the basic generators of the braid group are mapped into the Temperley–Lieb algebra. Every knot arises as the closure of a braid; the invariant arises by mapping the open braid into the Temperley–Lieb algebra and taking the trace there.

Moreover, it turns out that this connection can itself carry interesting information between the computer science ideas and the geometry and algebra. Indeed, using computer science methods, it is possible to give the first **direct presentation** (no quotients) of the Temperley–Lieb algebra, using logical methods. In fact, the elements of the Temperley–Lieb algebra are completely determined by the relations they induce on the “dots,” and **planarity** can be characterized using only the ordering relations on the two rows of dots. Moreover, the multiplication of the algebra can be described as a form of cut elimination, using the methods developed in the “geometry of interaction” (Abramsky 1996; Abramsky and Jagadeesan 1994a; Girard 1989). This exactly corresponds to characterizing the **geometric composition of diagrams** as above in terms of an **information flow dynamics**, just as in the case of multiplicative linear logic. This also shows that **planarity is an invariant of cut elimination** and raises interesting questions about computational expressiveness under topological constraints such as planarity and the computational significance of braiding.

Applicative Computation

These ideas in turn apply directly to **applicative computation**, offering the same combination of

geometric/diagrammatic and information-flow tools of analysis. This is illustrated firstly with a **planar combinator**

$$\mathbf{B} \equiv \lambda x. \lambda y. \lambda z. x(yz) : (B \rightarrow C) \rightarrow (A \rightarrow B) \rightarrow (A \rightarrow C)$$


characterized by the equation

$$\mathbf{B}abc = a(bc).$$


The interpretation of the open term

$$x : B \rightarrow C, y : A \rightarrow B, z : A \vdash x(yz) : C$$


is as follows:



Here x^+ is the output of x , and x^- the input, and similarly for y . The output of the whole expression is o . When the variables are abstracted, the following caps-only diagram is obtained:




Now consider an application $\mathbf{B}abc$:



If the constraint on planarity in the diagrams is relaxed, a similar representation can be given of the **commutation combinator**:

$$C \equiv \lambda x.\lambda y.\lambda z.xyz :$$
$$(A \rightarrow B \rightarrow C) \rightarrow B \rightarrow A \rightarrow C$$




More generally, the **Brauer algebra** (1931) arises if the planarity condition on the TL algebra is removed. This plays an important role in the representation theory of the orthogonal group (“Schur–Weyl duality”), which is now part of a whole genre of “diagram algebras” in representation theory.


With **BCI** combinators, one can interpret **linear λ -calculus**. The Kelly–Laplaza construction of the free compact closed category (Kelly and Laplaza 1980) can be retrieved by a straightforward generalization of these ideas.

Quantum Computation


The **same** graphical calculus and underlying algebraic structure which have been tracked through logic, geometry, and computation has been applied to quantum information and computation, yielding an incisive analysis of **quantum information flow**, and powerful and illuminating methods for reasoning about quantum informatic processes and protocols (Abramsky and Coecke 2004). This “strongly compact closed graphical calculus” can be seen as a very substantial two-dimensional extension of Dirac’s *bra-ket* notation (PAM 1947). In the graphical calculus, physical processes are depicted by boxes, and the inputs and outputs of these boxes are labeled by *types* which indicate the kind of system on which these boxes act, e.g., one qubit, several qubits, classical data, etc.: see Fig. 5. Algebraically, these correspond to $1_A: A \rightarrow A, f: A \rightarrow B, g \circ f, 1_A \otimes 1_B, f \otimes 1_C, f \otimes g, (f \otimes g) \circ h$, respectively. (The convention in these diagrams is that the “upward” vertical direction represents progress of time.)



Logic and Geometry of Agents in Agent-Based Modeling, Fig. 5 Graphical calculus for monoidal categories




Logic and Geometry of Agents in Agent-Based Modeling, Fig. 6 Linear Negation



Logic and Geometry of Agents in Agent-Based Modeling, Fig. 7 Linear Conjunction - no information flow

Logic and Geometry of Agents in Agent-Based Modeling, Fig. 8 Linear Disjunction - no information flow




Kets, Bras, and Scalars A special role is played by boxes with either no input or no output, corresponding to **states** and **costates**, respectively (cf. Dirac’s kets and bras (PAM 1947)), which are depicted by triangles. **Scalars** then arise naturally by composing these elements (cf. inner product or Dirac’s bra-ket):


$$\begin{array}{c} | \\ \downarrow \\ \psi \end{array} \quad \begin{array}{c} \pi \\ \uparrow \\ A \end{array} \quad \diamond \pi \circ \psi = \begin{array}{c} \pi \\ \hline \psi \end{array}$$

Bell States and Costates The cups and caps which have already appeared in their various


Logic and Geometry of Agents in Agent-Based Modeling, Fig. 9 The Copy-Cat Strategy



Logic and Geometry of Agents in Agent-Based Modeling, Fig. 10 Composition as Interaction




Logic and Geometry of Agents in Agent-Based Modeling, Fig. 11 Teleportation




guises as axiom and cut links, or in abstraction and application, now take on the role of **Bell states and costates**, the fundamental building blocks of

quantum entanglement. (Mathematically, they arise as the transpose and co-transpose of the identity, which exist in any finite-dimensional


Hilbert space by “map-state duality.”) They are represented with triangles encasing the cups and caps to emphasize their **operational character** in the physical interpretation: they represent **preparation and test** of the Bell state.



The formation of **names** and **conames** of arrows (i.e., map-state and map-costate duality) is conveniently depicted thus:




The key lemma in exposing the quantum information flow in (bipartite) entangled quantum systems (Abramsky and Coecke 2004):



Note in particular the interesting phenomenon of “apparent reversal of the causal order.” While on the left, physically, **the** state labeled g is firstly prepared, and then the costate labeled f is applied, the global effect is *as if* f itself is applied first and only then g .

Derivation of Quantum Teleportation This is the most basic application of compositionality in action. The basic quantum mechanical potential for teleportation can be read off immediately from the geometry of Bell states and costates:

This is not quite the whole story, because of the non-deterministic nature of measurements. But it suffices to introduce a unitary correction. Using the lemma, the full description of teleportation becomes:



Further Directions

This article has described what is still very much an emerging area of research, rather than surveying an established field. Based on the progress which has already been made, a number of promising directions for future work are apparent, which may lead to important contributions to the general study of agent-based systems:

- The way in which, as illustrated in the previous subsection, the **same** structures arise in a wide diversity of situations strongly suggests that a logic of interaction can be developed, which combines genuine depth with wide applicability.
- Further study of the links between logic and geometry, and of the significance of geometric constraints on informational contexts, looks promising.
- The program of full completeness, and giving equivalent geometric and interactive characterizations of constrained type systems which capture important complexity classes, may open up new insights into currently intractable problems.
- The connections between informatic and physical structures and constraints which have been found in the work on quantum information are promising both as new directions in the foundations of quantum mechanics and the quantum/classical boundary and in elucidating fundamental structures of informatics.

Bibliography

Primary Literature

Abramsky S (1996) Retracing some paths in process algebra. In: Proceedings of CONCUR 96. Lectures notes in computer science, vol 1119. Springer, Berlin, pp 1–17

- Abramsky S (2002) Algorithmic game semantics: a tutorial introduction. In: *Proof and system-reliability*. Kluwer, Dordrecht
- Abramsky S (2004) High-level methods for quantum computation and information. In: *Proceedings of the 19th annual IEEE symposium on logic in computer science*. IEEE Computer Science Press, Los Alamitos, pp 410–414
- Abramsky S (2005) Abstract scalars, loops, and free traced and strongly compact closed categories. In: *Proceedings of CALCO 2005*. Springer lecture notes in computer science, vol 3629. Springer, Berlin, pp 1–31
- Abramsky S (2007) Temperley–Lieb algebras and geometry of interaction. In: Chen G, Kauffman L, Lomonaco S (eds) *Mathematics of quantum computing and technology*. Chapman and Hall/CRC, Boca Raton, pp 413–458
- Abramsky S, Coecke B (2004) A categorical semantics of quantum protocols. In: *Proceedings of the 19th Annual IEEE symposium on logic in computer science: LICS 2004*. IEEE Computer Society, Los Alamitos, pp 415–425
- Abramsky S, Coecke B (2005) Abstract physical traces. *Theory Appl Categ* 14:111–124
- Abramsky S, Jagadeesan R (1994a) New foundations for the geometry of interaction. *Inf Comput* 111:53–119
- Abramsky S, Jagadeesan R (1994b) Games and full completeness for multiplicative linear logic. *J Symb Log* 59:543–574
- Abramsky S, Jagadeesan R (2005) A game semantics for generic polymorphism. *Ann Pure Appl Log* 133:3–37
- Abramsky S, McCusker G (1997) Linearity, sharing and state. In: O’Hearn P, Tennent RD (eds) *Algol-like languages*. Birkhauser, Basel, pp 317–348
- Abramsky S, McCusker G (1999a) Full abstraction for idealized Algol with passive expressions. *Theor Comput Sci* 227:3–42
- Abramsky S, McCusker G (1999b) Game semantics. In: *Computational logic: Proceedings of the 1997 Marktoberdorf Summer School*. Springer, Berlin, pp 1–56
- Abramsky S, Mellies P-A (1999) Concurrent games and full completeness. In: *Proceedings of the 14th international symposium on logic in computer science*. Computer Society Press of the IEEE, Los Alamitos, pp 431–442
- Abramsky S, Honda K, McCusker G (1998) A fully abstract game semantics for general references. In: *Proceedings of the 13th international symposium on logic in computer science*. Computer Society Press of the IEEE, Los Alamitos, pp 334–344
- Abramsky S, Jagadeesan R, Malacaria P (2000) Full abstraction for PCF. *Inf Comput* 163:409–470
- Abramsky S, Ghica DR, Murawski AS, Ong C-HL (2004a) Applying game semantics to compositional software modeling and verification. In: *Proceedings of the TACAS’04*. LNCS, vol 2988. pp 421–435
- Abramsky S, Ghica DR, Murawski AS, Stark IDB, Ong C-HL (2004b) Nominal games and full abstraction for the nu-calculus. In: *Proceedings of the LICS’04*. IEEE Computer Society Press, Los Alamitos, pp 150–159
- Blute R, Scott PJ (1998) The Shuffle Hopf Algebra and noncommutative full completeness. *J Symb Log* 63(4):1413–1436
- Blute R, Hamano M, Scott PJ (2005) Softness of hypercoherences and MALL full completeness. *Ann Pure Appl Log* 131(1–3):1–63
- Danos V, Harmer R (2002) Probabilistic game semantics. *ACM Trans Comput Log* 3(3):359–382
- Devarajan H, Hughes D, Plotkin G, Pratt V (1999) Full completeness of the multiplicative linear logic of Chu spaces. In: *Proceedings of the 14th Annual IEEE symposium on logic in computer science*. pp 234–242
- Ghica DR, McCusker G (2000) Reasoning about idealized algol using regular languages. In: *Proceedings of the ICALP’00*. LNCS, vol 1853. pp 103–116
- Ghica DR, Murawski AS (2004) Angelic semantics of finegrained concurrency. In: *Proceedings of the FOSSACS’04*. LNCS, vol 2987. pp 211–225
- Ghica DR, Murawski AS (2006) Compositional model extraction for higher-order concurrent programs. In: *Proceedings of the TACAS’06*. LNCS
- Girard J-Y (1987) Linear logic. *Theor Comput Sci* 50(1):1–102
- Girard J-Y (1989) Geometry of interaction I: interpretation of system F. In: Ferro R et al (eds) *Logic Colloquium ‘88*. Elsevier, Amsterdam, pp 221–260
- Harmer R, McCusker G (1999) A fully abstract game semantics for finite nondeterminism. In: *Proceedings of the 14th Annual IEEE symposium on logic in computer science*. IEEE Computer Society Press, Los Alamitos
- Hughes D (2000) *Hypergame semantics: full completeness for system F*. D Phil mathematical sciences. Oxford University, Oxford
- Hyland JME, Ong C-HL (2000) On full abstraction for PCF.i. Models, observables and the full abstraction problem, ii. Dialogue games and innocent strategies, iii. A fully abstract and universal game model. *Inf Comput* 163:285–408
- Kauffman LH (1990) An invariant of regular isotopy. *Trans Am Math Soc* 318(2):417–471
- Kelly GM, Laplaza ML (1980) Coherence for compact closed categories. *J Pure Appl Algebra* 19:193–213
- Laird J (1997) Full abstraction for functional languages with control. In: *Proceedings of the 12th Annual symposium on logic in computer science, LICS ‘97*. Extended abstract
- Laird J (2001) A fully abstract games semantics of local exceptions. In: *Proceedings of the 16th Annual symposium on logic in computer science, LICS ‘01*. Extended abstract
- Loader R (1994) *Models of Lambda calculi and linear logic*. PhD thesis, Oxford University, Oxford
- Murawski AS, Ong C-HL, Walukiewicz I (2005) Idealized Algol with ground recursion and DPDA equivalence. In: *Proceedings of the ICALP’05*. LNCS, vol 3580. pp 917–929

PAM D (1947) Principles of quantum mechanics. Oxford University Press, Oxford

VFR J (1985) A polynomial invariant for links via von Neumann algebras. Bull Am Math Soc 129:103–112

Books and Reviews

Girard J-Y, Lafont Y, Taylor P (1989) Proof and types. Cambridge Tracts in Theoretical Computer Science

Hindley JR, Seldin JP (1986) Introduction to combinators and the λ -calculus. Cambridge University Press, Cambridge

Kauffman LH (1994) Knots in physics. World Scientific Press, Singapore

Troelstra AS (1992) Lectures on linear logic. Center for the Study of Language and Information Lecture Notes No. 29

Agent-Based Modeling and Artificial Life

Charles M. Macal

Center for Complex Adaptive Agent Systems
Simulation (CAS2), Decision and Information
Sciences Division, Argonne National Laboratory,
Argonne, IL, USA

Article Outline

Glossary
Definition of the Subject
Introduction
Artificial Life
ALife in Agent-Based Modeling
Future Directions
Bibliography

Glossary

Adaptation The process by which organisms (agents) change their behavior or by which populations of agents change their collective behaviors with respect to their environment.

Agent-based modeling (ABM) A modeling and simulation approach applied to a complex system or complex adaptive system, in which the model is comprised of a large number of interacting elements (agents).

Ant colony optimization (ACO) A heuristic optimization technique motivated by collective decision processes followed by ants in foraging for food.

Artificial chemistry Chemistry based on the information content and transformation possibilities of molecules.

Artificial life (ALife) A field that investigates life's essential qualities primarily from an information content perspective.

Artificial neural network (ANN) A heuristic optimization and simulated learning technique

motivated by the neuronal structure of the brain.

Autocatalytic set A closed set of chemical reactions that is self-sustaining.

Autonomous The characteristic of being capable of making independent decisions over a range of situations.

Avida An advanced artificial life computer program developed by Adami (Adami 1998) and others that models populations of artificial organisms and the essential features of life such as interaction and replication.

Biologically inspired computational algorithm Any kind of algorithm that is based on biological metaphors or analogies.

Cellular automaton (CA) A mathematical construct and technique that models a system in discrete time and discrete space in which the state of a cell depends on transition rules and the states of neighboring cells.

Coevolution A process by which many entities adapt and evolve their behaviors as a result of mutually effective interactions.

Complex system A system comprised of a large number of strongly interacting components (agents).

Complex adaptive system (CAS) A system comprised of a large number of strongly interacting components (agents) that adapt at the individual (agent) level or collectively at the population level.

Decentralized control A feature of a system in which the control mechanisms are distributed over multiple parts of the system.

Digital organism An entity that is represented by its essential information-theoretic elements (genomes) and implemented as a computational algorithm or model.

Downward causation The process by which a higher-order emergent structure takes on its own emergent behaviors and these behaviors exert influence on the constituent agents of the emergent structure.

- Dynamic network analysis** Network modeling and analysis in which the structure of the network, i.e., nodes (agents) and links (agent interactions), is endogenous to the model.
- Echo** An artificial life computer program developed by Holland (Holland 1975) that models populations of complex adaptive systems and the essential features of adaptation in nature.
- Emergence** The process by which order is produced in nature.
- Entropy** A measure of order, related to the information needed to specify a system and its state.
- Evolution (artificial)** The process by which a set of instructions is transmitted and changes over successive generations.
- Evolutionary game** A repeated game in which agents adapt their strategies in recognition of the fact that they will face their opponents in the future.
- Evolution strategies** A heuristic optimization technique motivated by the genetic operation of selection and mutation.
- Evolutionary algorithm** Any algorithm motivated by the genetic operations including selection, mutation, crossover, etc.
- Evolutionary computing** A field of computing based on the use of evolutionary algorithms.
- Finite state machine** A mathematical model consisting of entities with a finite (usually small) number of possible discrete states.
- Game of Life, Life** A cellular automaton developed by Conway (Berlekamp et al. 2003) that illustrates a maximally complex system based on simple rules.
- Generative social science** Social science investigation with the goal of understanding how social processes emerge out of social interaction.
- Genetic algorithm (GA)** A specific kind of evolutionary algorithm motivated by the genetic operations of selection, mutation, and crossover.
- Genetic programming (GP)** A specific kind of evolutionary algorithm that manipulates symbols according to prescribed rules, motivated by the genetic operations of selection, mutation, and crossover.
- Genotype** A set of instructions for a developmental process that creates a complex structure, as in a genotype for transmitting genetic information and seeding a developmental process leading to a phenotype.
- Hypercycle** A closed set of functional relations that is self-sustaining, as in an autocatalytic chemical reaction network.
- Individual-based model** An approach originating in ecology to model populations of agents that emphasizes the need to represent diversity among individuals.
- Langton's ant** An example of a very simple computational program that computes patterns of arbitrary complexity after an initial series of simple structures.
- Langton's loop** An example of a very simple computational program that computes replicas of its structures according to simple rules applied locally as in cellular automata.
- Learning classifier system (LCS)** A specific algorithmic framework for implementing an adaptive system by varying the weights applied to behavioral rules specified for individual agents.
- Lindenmeyer system (L-system)** A formal grammar, which is a set of rules for rewriting strings of symbols.
- Machine learning** A field of inquiry consisting of algorithms for recognizing patterns in data (e.g., data mining) through various computerized learning techniques.
- Mind-body problem** A field of inquiry that addresses how human consciousness arises out of material processes and whether consciousness is the result of a logical-deductive or algorithmic process.
- Meme** A term coined by Dawkins (Dawkins 1989) to refer to the minimal encoding of cultural information, similar to the genome's role in transmitting genetic information.
- Particle swarm optimization** An optimization technique similar to ant colony optimization, based on independent particles (agents) that search a landscape to optimize a single objective or goal.
- Phenotype** The result of an instance of a genotype interacting with its environment through a developmental process.
- Reaction-diffusion system** A system that includes mechanisms for both attraction and

transformation (e.g., of agents) as well as repulsion and diffusion.

Recursively generated object An object that is generated by the repeated application of simple rules.

Self-organization A process by which structure and organization arise from within the endogenous instructions and processes inherent in an entity.

Self-replication The process by which an agent (e.g., organism, machine, etc.) creates a copy of itself that contains instructions for both the agent's operation and its replication.

Social agent-based modeling Agent-based modeling applied to social systems, generally applied to people and human populations, but also animals.

Social network analysis (SNA) A collection of techniques and approaches for analyzing networks of social relationships.

Stigmergy The practice of agents using the environment as a tool for communication with other agents to supplement direct agent-to-agent communication.

Swarm An early agent-based modeling toolkit designed to model artificial life applications.

Swarm intelligence Collective intelligence based on the actions of a set of interacting agents behaving according to a set of prescribed simple rules.

Symbolic processing A computational technique that consists of processing symbols rather than strictly numerical data.

Sugarscape An abstract agent-based model of artificial societies developed by Epstein and Axtell (Epstein and Axtell 1996) to investigate the emergence of social processes.

Tierra An early artificial life computer program developed by Ray (Ray 1991) that models populations of artificial organisms and the essential features of life such as interaction and replication.

Universal Turing machine (UTM) An abstract representation of the capabilities of any computable system.

Update rule A rule or transformation directive for changing or updating the state of an entity or agent, as, for example, updating the state of

an agent in an agent-based model or updating the state of an L-system.

Definition of the Subject

Agent-based modeling began as the computational arm of artificial life some 20 years ago. Artificial life is concerned with the emergence of order in nature. How do systems self-organize themselves and spontaneously achieve a higher-ordered state? Agent-based modeling, then, is concerned with exploring and understanding the processes that lead to the emergence of order through computational means. The essential features of artificial life models are translated into computational algorithms through agent-based modeling. With its historical roots in artificial life, agent-based modeling has become a distinctive form of modeling and simulation. Agent-based modeling is a bottom-up approach to modeling complex systems by explicitly representing the behaviors of large numbers of agents and the processes by which they interact. These essential features are all that is needed to produce at least rudimentary forms of emergent behavior at the systems level. To understand the current state of agent-based modeling and where the field aspires to be in the future, it is necessary to understand the origins of agent-based modeling in artificial life.

Introduction

The field of artificial life, or “ALife,” is intimately connected to agent-based modeling, or “ABM.” Although one can easily enumerate some of life's distinctive properties, such as reproduction, respiration, adaptation, emergence, etc., a precise definition of life remains elusive.

Artificial life had its inception as a coherent and sustainable field of investigation at a workshop in the late 1980s (Langton 1989a). This workshop drew together specialists from diverse fields who had been working on related problems in different guises, using different vocabularies suited to their fields.

At about the same time, the introduction of the personal computer suddenly made computing accessible, convenient, inexpensive, and compelling as an experimental tool. The future seemed to have almost unlimited possibilities for the development of ALife computer programs to explore life and its possibilities. Thus, several ALife software programs emerged that sought to encapsulate the essential elements of life through incorporation of ALife-related algorithms into easily usable software packages that could be widely distributed. Computational programs for modeling populations of digital organisms, such as *Tierra*, *Avida*, and *Echo*, were developed along with more general purpose agent-based simulators such as *Swarm*.

Yet, the purpose of ALife was never restricted to understanding or recreating life as it exists today. According to Langton:

Artificial systems which exhibit lifelike behaviors are worthy of investigation on their own right, whether or not we think that the processes that they mimic have played a role in the development or mechanics of life as we know it to be. Such systems can help us expand our understanding of life as it could be. (p. xvi in Langton 1989a)

The field of ALife addresses lifelike properties of systems at an abstract level by focusing on the information content of such systems independent of the medium in which they exist, whether it be biological, chemical, physical, or in silico. This means that computation, modeling, and simulation play a central role in ALife investigations.

The relationship between ALife and ABM is complex. A case can be made that the emergence of ALife as a field was essential to the creation of agent-based modeling. Computational tools were both required and became possible in the 1980s for developing sophisticated models of digital organisms and general purpose artificial life simulators. Likewise, a case can be made that the possibility for creating agent-based models was essential to making ALife a promising and productive endeavor. ABM made it possible to understand the logical outcomes and implications of ALife models and lifelike processes. Traditional analytical means, although valuable in establishing baseline information, were limited in their capabilities to include essential features of ALife. Many threads of ALife are still

intertwined with developments in ABM and vice versa. Agent-based models demonstrate the emergence of lifelike features using ALife frameworks; ALife algorithms are widely used in agent-based models to represent agent behaviors. These threads are explored in this entry. In ALife terminology, one could say that ALife and ABM have coevolved to their present states. In all likelihood, they will continue to do so.

This entry covers in a necessarily brief and perhaps superficial, but broad, way these relationships between ABM and ALife and extrapolates to future possibilities. This entry is organized as follows. Section “[Artificial Life](#)” introduces artificial life, its essential elements, and its relationship to computing and agent-based modeling. Section “[ALife in Agent-Based Modeling](#)” describes several examples of ABM applications spanning many scales. Section “[Future Directions](#)” concludes with future directions for ABM and ALife. A bibliography is included for further reading.

Artificial Life

Artificial life was initially motivated by the need to model biological systems and brought with it the need for computation. The field of ALife has always been multidisciplinary and continues to encompass a broad research agenda covering a variety of topics from a number of disciplines, including:

- Essential elements of life and artificial life
- Origins of life and self-organization
- Evolutionary dynamics
- Replication and development processes
- Learning and evolution
- Emergence
- Computation of living systems
- Simulation systems for studying ALife
- Many others

Each of these topics has threads leading into agent-based modeling.

The Essence of ALife

The essence of artificial life is summed up by Langton (p. xxii in Langton (1989a)) with a list of essential characteristics:

- Lifelike behavior on the part of man-made systems
- Semiautonomous entities whose local interactions with one another are governed by a set of simple rules
- Populations, rather than individuals
- Simple rather than complex specifications
- Local rather than global control
- Bottom-up rather than top-down modeling
- Emergent rather than prespecified behaviors


Langton observes that complex high-level dynamics and structures often emerge (in living and artificial systems), developing over time out of the local interactions among low-level primitives. Agent-based modeling has grown up around the need to model the essentials of ALife.

Self-Replication and Cellular Automata

Artificial life traces its beginnings to the work of John von Neumann in the 1940s and investigations into the theoretical possibilities for developing a *self-replicating* machine (Taub 1961). Such a self-replicating machine carries instructions not only for its operations but also for its replication. The issue is concerned with how to replicate such a machine that contained the instructions for its operation along with the instructions for its

replication. Did a machine need to contain both the instructions for the machine's operation and replication, as well as instructions for replicating the instructions on how to replicate the original machine? (see Fig. 1.) Von Neumann used the abstract mathematical construct of *cellular automata*, originally conceived in discussions with Stanislaw Ulam, to prove that such a machine could be designed, at least in theory. Von Neumann was never able to build such a machine due to the lack of sophisticated computers that existed at the time.

Cellular automata (CA) have been central to the development of computing artificial life models. Virtually all of the early agent-based models that required agents to be spatially located were in the form of von Neumann's original cellular automata. A cellular automaton is a *finite-state machine* in which time and space are treated as discrete rather than continuous, as would be the case, for example, in differential equation models. A typical CA is a two-dimensional grid or lattice consisting of cells. Each cell assumes one of a finite number of states at any time. A cell's neighborhood is the set of cells surrounding a cell, typically a five-cell neighborhood (von Neumann neighborhood) or a nine-cell neighborhood (Moore neighborhood), as in Fig. 2.



Agent-Based Modeling and Artificial Life, Fig. 1 Von Neumann's self-replication problem

Agent-Based Modeling and Artificial Life,
Fig. 2 Cellular automata neighborhoods



A set of simple state transition rules determines the value of each cell based on the cell’s state and the states of neighboring cells. Every cell is updated at each time according to the transition rules. Each cell is identical in terms of its update rules. Cells differ only in their initial states. A CA is deterministic in the sense that the same state for a cell and its set of neighbors always results in the same updated state for the cell. Typically, CAs are set up with periodic boundary conditions, meaning that the set of cells on one edge of the grid boundary is the neighbor cells to the cells on the opposite edge of the grid boundary. The space of the CA grid forms a surface on a toroid, or donut-shape, so there is no boundary per se. It is straightforward to extend the notion of cellular automata to two, three, or more dimensions.

Von Neumann solved the self-replication problem by developing a cellular automaton in which each cell had 29 possible states and five neighbors (including the updated cell itself). In the von Neumann neighborhood, neighbor cells are in the north, south, east, and west directions from the updated cell.

The Game of Life Conway’s Game of Life, or Life, developed in the 1970s, is an important example of a CA (Berlekamp et al. 2003; Gardner 1970; Poundstone 1985). The simplest way to illustrate some of the basic ideas of agent-based modeling is through a CA. The Game of Life is a two-state, nine-neighbor cellular automaton with three rules that determine the state (either On, i.e., shaded, or Off, i.e., white) of each cell:


1. A cell will be On in the next generation if exactly three of its eight neighboring cells are currently On.
2. A cell will retain its current state if exactly two of its neighbors are On.
3. A cell will be Off otherwise.

Initially, a small set of On cells is randomly distributed over the grid. The three rules are then applied repeatedly to all cells in the grid.

After several updates of all cells on the grid, distinctive patterns emerge, and in some cases these patterns can sustain themselves indefinitely throughout the simulation (Fig. 3). The state of each cell is based only on the current state of the cell and the cells touching it in its immediate neighborhood. The nine-neighbor per neighborhood assumption built into Life determines the scope of the locally available information for each cell to update its state.

Conway showed that, at least in theory, the structures and patterns that can result during a Life computation are complex enough to be the basis for a fully functional computer that is complex enough to spontaneously generate self-replicating structures (see the section below on universal computation). Two observations are important about the Life rules:

- As simple as the state transition rules are, by using only local information, structures of arbitrarily high complexity can emerge in a CA.
- The specific patterns that emerge are extremely sensitive to the specific rules used. For example, changing Rule 1 above to “A cell will be



Agent-Based Modeling and Artificial Life, Fig. 3 Game of life

On in the next generation if exactly four of its eight neighboring cells are currently On” results in the development of completely different patterns.

- The Game of Life provides insights into the role of information in fundamental life processes.

Cellular Automata Classes Wolfram investigated the possibilities for complexity in cellular automata across the full range of transition rules and initial states, using one-dimensional cellular automata (Wolfram 1984). He categorized four distinct classes for the resulting patterns produced by a CA as it is solved repeatedly over time. These are:

- Class I: homogeneous state
- Class II: simple stable or periodic structure
- Class III: chaotic (non-repeating) pattern
- Class IV: complex patterns of localized structures

The most interesting of these is Class IV cellular automata, in which very complex patterns of non-repeating localized structures emerge that are often long lived. Wolfram showed that these Class IV structures were also complex enough to support universal computation (Wolfram 2002). Langton (1992) coined the term “life at the edge of chaos” to describe the idea that Class IV systems are situated in a thin region between Class II and Class III systems. Agent-based models often yield Class I, Class II, and Class III behaviors.

Other experiments with CAs investigated the simplest representations that could replicate

themselves and produce emergent structures. *Langton’s loop* is a self-replicating two-dimensional cellular automaton, much simpler than von Neumann’s (Langton 1984). Although not complex enough to be a universal computer, Langton’s loop was the simplest known structure that could reproduce itself. *Langton’s ant* is a two-dimensional CA with a simple set of rules, but complicated emergent behavior. Following a simple set of rules for moving from cell to cell, a simulated ant displays unexpectedly complex behavior. After an initial period of chaotic movements in the vicinity of its initial location, the ant begins to build a recurrent pattern of regular structures that repeats itself indefinitely (Langton 1986). Langton’s ant has behaviors complex enough to be a universal computer.


Genotype/Phenotype Distinction

Biologists distinguish between the *genotype* and the *phenotype* as hallmarks of biological systems. The genotype is the template – the set of instructions, the specification, and the blueprint – for an organism. DNA is the genotype for living organisms, for example. A DNA strand contains the complete instructions for the replication and development of the organism. The phenotype is the organism – the machine, the product, and the result – that develops from the instructions in the genotype (Fig. 4).

Morphogenesis is the developmental process by which the phenotype develops in accord with the genotype, through interactions with and resources

Agent-Based Modeling and Artificial Life,

Fig. 4 Genotype and phenotype relations



obtained from its environment. In a famous paper, Turing (1952) modeled the dynamics of morphogenesis and, more generally, the problem of how patterns self-organize spontaneously in nature. Turing used differential equations to model a simple set of reaction-diffusion chemical reactions. Turing demonstrated that only a few assumptions were necessary to bring about the emergence of wave patterns and gradients of chemical concentration, suggestive of morphological patterns that commonly occur in nature. *Reaction-diffusion systems* are characterized by the simultaneous processes of attraction and repulsion and are the basis for the agent behavioral rules (attraction and repulsion) in many social agent-based models.

More recently, Bonabeau extended Turing's treatment of morphogenesis to a theory of pattern formation based on agent-based modeling. Bonabeau (1997) states the reason for relying on ABM: "because pattern-forming systems based on agents are (relatively) more easily amenable to experimental observations."

Information Processes One approach to building systems from a genotype specification is based on the methodology of *recursively generated objects*. Such recursive systems are compact in their specification, and their repeated application can result in complex structures, as demonstrated by cellular automata.

Recursive systems are logic systems in which strings of symbols are recursively rewritten based on a minimum set of instructions. Recursive systems, or term replacement systems, as they have been called, can result in complex structures. Examples of recursive systems include cellular automata, as described above, and *Lindenmayer systems*, called *L-systems* (Le Novere and Shimizu 2001). An L-system consists of a formal grammar, which is a set of rules for rewriting strings of symbols. L-systems have been used extensively for modeling living systems, for example, plant growth and development, producing highly realistic renderings of plants, with intricate morphologies and branching structures.

Wolfram (1999) used symbolic recursion as a basis for developing *Mathematica*, the computational mathematics system based on *symbolic processing* and term replacement. Unlike numeric programming languages, a symbolic programming language allows a variable to be a basic object and does not require a variable to be assigned a value before it is used in a program.

Any agent-based model is essentially a recursive system. Time is simulated by the repeated application of the agent updating rules. The genotype is the set of rules for the agent behaviors. The phenotype is a set of the patterns and structures that emerge from the computation. As in cellular automata and recursive systems, extremely complex structures emerge in agent-based models that are often unpredictable from examination of the agent rules.

Emergence One of the primary motivations for the field of ALife is to understand *emergent* processes, that is, the processes by which life emerges from its constituent elements. Langton writes: “The ‘key’ concept in ALife, is emergent behavior” (p. 2 in Langton 1989b). Complex systems exhibit patterns of emergence that are not predictable from inspection of the individual elements. Emergence is described as unexpected, unpredictable, or otherwise surprising. That is, the modeled system exhibits behaviors that are not explicitly built into the model. Unpredictability is due to the nonlinear effects that result from the interactions of entities having simple behaviors. Emergence by these definitions is something of a subjective process.

In biological systems, emergence is a central issue whether it be the emergence of the phenotype from the genotype, the emergence of protein complexes from genomic information networks (Kauffman 1993), or the emergence of consciousness from networks of millions of brain cells.

One of the motivations for agent-based modeling is to explore the emergent behaviors exhibited by the simulated system. In general, agent-based models often exhibit patterns and relationships that emerge from agent interactions. An example is the observed formation of groups of agents that collectively act in coherent and coordinated patterns. Complex adaptive systems, widely investigated by Holland in his agent-based model Echo (Holland 1995), are often structured in hierarchies of emergent structures. Emergent structures can collectively form higher-order structures, using the lower-level structures as building blocks. An emergent structure itself can take on new emergent behaviors. These structures in turn affect the agents from which the structure has emerged in a process called *downward causation* (Gilbert 2002). For example, in the real world, people organize and identify with groups, institutions, nations, etc. They create norms, laws, and protocols that in turn act on the individuals comprising the group.

From the perspective of agent-based modeling, emergence has some interesting challenges for modeling:

- How does one operationally define emergence with respect to agent-based modeling?

- How does one automatically identify and measure the emergence of entities in a model?
- How do agents that comprise an emergent entity perceived by an observer recognize that they are part of that entity?

Artificial Chemistry

Artificial chemistry is a subfield of ALife. One of the original goals of artificial chemistry was to understand how life could originate from prebiotic chemical processes. Artificial chemistry studies self-organization in chemical reaction networks by simulating chemical reactions between artificial molecules. Artificial chemistry specifies well-understood chemical reactions and other information such as reaction rates, relative molecular concentrations, probabilities of reaction, etc. These form a network of possibilities. The artificial substances and the networks of chemical reactions that emerge from the possibilities are studied through computation. Reactions are specified as recursive algebras and activated as term replacement systems (Fontana 1992).

Hypercycles The emergence of *autocatalytic sets*, or *hypercycles*, has been a prime focus of artificial chemistry (Eigen and Schuster 1979). A hypercycle is a self-contained system of molecules and a self-replicating, and thereby self-sustaining, cyclic linkage of chemical reactions. Hypercycles evolve through a process by which self-replicating entities compete for selection.

The hypercycle model illustrates how an ALife process can be adopted to the agent-based modeling domain. Inspired by the hypercycle model, Padgett et al. (2003) developed an agent-based model of the coevolution of economic production and economic firms, focusing on skills. Padgett used the model to establish three principles of social organization that provide foundations for the evolution of technological complexity:

- Structured topology (how interaction networks form)
- Altruistic learning (how cooperation and exchange emerge)

- Stigmergy (how agent communication is facilitated by using the environment as a means of information exchange among agents)

Digital Organisms The widespread availability of personal computers spurred the development of ALife programs used to study evolutionary processes in silico. Tierra was the first system devised in which computer programs were successfully able to evolve and adapt (Ray 1991). Avida extended Tierra to account for the spatial distribution of organisms and other features (Ofria and Wilke 2004; Wilke and Adami 2002). Echo is a simulation framework for implementing models to investigate mechanisms that regulate diversity and information processing in complex adaptive systems (CAS), systems comprised of many interacting adaptive agents (Holland 1975, 1995). In implementations of Echo, populations evolve interaction networks, resembling species communities in ecological systems, which regulate the flow of resources.

Systems such as Tierra, Avida, and Echo simulate populations of digital organisms, based on the genotype/phenotype schema. They employ computational algorithms to mutate and evolve populations of organisms living in a simulated computer environment. Organisms are represented as strings of symbols, or agent attributes, in computer memory. The environment provides them with resources (computation time) they need to survive, compete, and reproduce. Digital organisms interact in various ways and develop strategies to ensure survival in resource-limited environments.

Digital organisms are extended to agent-based modeling by implementing individual-based models of food webs in a system called DOVE (Wilke and Chow 2006). Agent-based models allow a more complete representation of agent behaviors and their evolutionary adaptability at both the individual and population levels.

ALife and Computing

Creating lifelike forms through computation is central to artificial life. Is it possible to create life through computation? The capabilities and limitations of computation constrain the types of artificial life that can be created. The history of ALife

has close ties with important events in the history of computation.

Alan Turing (1938) investigated the limitations of computation by developing an abstract and idealized computer, called a *universal Turing machine* (UTM). A UTM has an infinite tape (memory) and is therefore an idealization of any actual computer that may be realized. A UTM is capable of computing anything that is computable, that is, anything that can be derived via a logical, deductive series of statements. Are the algorithms used in today's computers, and in ALife calculations and agent-based models in particular, as powerful as universal computers?

Any system that can effectively simulate a small set of logical operations (such as AND and NOT) can effectively produce any possible computation. Simple rule systems in cellular automata were shown to be equivalent to universal computers (von Neumann 1966; Wolfram 2002) and in principal able to compute anything that is computable – *perhaps, even life!*

Some have argued that life, in particular human consciousness, is not the result of a logical-deductive or algorithmic process and therefore not computable by a universal Turing machine. This problem is more generally referred to as the *mind-body problem* (Lucas 1961). Dreyfus (1979) argues against the assumption often made in the field of artificial intelligence that human minds function like general purpose symbol manipulation machines. Penrose (1989) argues that the rational processes of the human mind transcend formal logic systems. In a somewhat different view, biological naturalism contends (Searle 1990) that human behavior might be able to be simulated, but human consciousness is outside the bounds of computation.

Such philosophical debates are as relevant to agent-based modeling as they are to artificial intelligence, for they are the basis of answering the question of what kind of systems and processes agent-based models will ultimately be able, or unable, to simulate.

Artificial Life Algorithms

ALife use several *biologically inspired computational algorithms* (Olariu and Zomaya 2006).

Bioinspired algorithms include those based on Darwinian evolution, such as evolutionary algorithms; those based on neural structures, such as neural networks; and those based on decentralized decision-making behaviors observed in nature. These algorithms are commonly used to model adaptation and learning in agent-based modeling or to optimize the behaviors of whole systems.

Evolutionary Computing

Evolutionary computing includes a family of related algorithms and programming solution techniques inspired by evolutionary processes, especially the genetic processes of DNA replication and cell division (Eiben and Smith 2007). These techniques are known as *evolutionary algorithms* and include the following (Back 1996):

- Genetic algorithms (Goldberg 1989, 1994; Holland 1975; Holland et al. 2000; Mitchell and Forrest 1994)
- Evolution strategies (Rechenberg 1973)
- Learning classifier systems (Holland et al. 2000).
- Genetic programming (Koza 1992)
- Evolutionary programming (Fogel et al. 1966).

Genetic algorithms (GA) model the dynamic processes by which populations of individuals evolve to improved levels of fitness for their particular environment over repeated generations. GAs illustrate how evolutionary algorithms process a population and apply the genetic operations of mutation and crossover (see Fig. 5). Each behavior is represented as a chromosome consisting of a series of symbols, for example, as a series of 0s and 1s. The encoding process establishing correspondence between behaviors and their chromosomal representations is part of the modeling process.

The general steps in a genetic algorithm are as follows:

1. *Initialization*: Generate an initial population of individuals. The individuals are unique and include specific encoding of attributes in chromosomes that represents the characteristics of the individuals.

2. *Evaluation*: Calculate the fitness of all individuals according to a specified fitness function.
3. *Checking*: If any of the individuals has achieved an acceptable level of fitness, stop; the problem is solved. Otherwise, continue with selection.
4. *Selection*: Select the best pair of individuals in the population for reproduction according to their high fitness levels.
5. *Crossover*: Combine the chromosomes for the two best individuals through a crossover operation, and produce a pair of offspring.
6. *Mutation*: Randomly mutate the chromosomes for the offspring.
7. *Replacement*: Replace the least fit individuals in the population with the offspring.
8. Continue at Step 2.


Steps 5 and 6 above, the operations of crossover and mutation, comprise the set of genetic operators inspired by nature. This series of steps for a GA comprise a basic framework rather than a specific implementation. Actual GA implementations include numerous variations and alternative implementations in several of the GA steps.

Evolution strategies (ES) are similar to genetic algorithms but rely on mutation as its primary genetic operator.

Learning classifier systems (LCS) build on genetic algorithms and adaptively assign relative weights to sensor-action sets that result in the most positive outcomes relative to a goal.

Genetic programming (GP) has similar features to genetic algorithms, but instead of using 0s and 1s or other symbols for comprising chromosomes, GPs combine logical operations and directives in a tree structure. In effect, chromosomes in GPs represent whole computer programs that perform a variety of functions with varying degrees of success and efficiencies. GP chromosomes are evaluated against fitness or performance measures and recombined. Better-performing chromosomes are maintained and expand their representation in the population. For example, an application of a GP is to evolve a better-performing rule set that represents an agent's behavior.

**Agent-Based Modeling
and Artificial Life,**
Fig. 5 Genetic algorithm



Evolutionary programming (EP) is a similar technique to genetic programming, but relies on mutation as its primary genetic operator.

Biologically Inspired Computing

Artificial neural networks (ANN) are another type of commonly used biologically inspired algorithm (Mehrotra et al. 1996). An artificial neural network uses mathematical models based on the structures observed in neural systems. An artificial neuron contains a stimulus-response model of neuron activation based on thresholds of stimulation. In modeling terms, neural networks are equivalent to nonlinear, statistical data modeling techniques. Artificial neural networks can be used to model complex relationships between inputs and outputs and to find patterns in data that are

dynamically changing. An ANN is adaptive in that changes in its structure are based on external or internal information that flows through the network. The adaptive capability makes ANN an important technique in agent-based models.

Swarm intelligence refers to problem-solving techniques, usually applied to solving optimization problems that are based on decentralized problem-solving strategies that have been observed in nature. These include:

- Ant colony optimization (Dorigo and Stützle 2004)
- Particle swarm optimization (Clerc 2006)

Swarm intelligence algorithms simulate the movement and interactions of large numbers of

ants or particles over a search space. In terms of agent-based modeling, the ants or particles are the agents, and the search space is the environment. Agents have position and state as attributes. In the case of particle swarm optimization, agents also have velocity.

Ant colony optimization (ACO) mimics techniques that ants use to forage and find food efficiently (Bonabeau et al. 1999; Engelbrecht 2006). The general idea of ant colony optimization algorithms is as follows:

1. In a typical ant colony, ants search randomly until one of them finds food.
2. Then they return to their colony and lay down a chemical pheromone trail along the way.
3. When other ants find such a pheromone trail, they are more likely to follow the trail rather than to continue to search randomly.
4. As other ants find the same food source, they return to the nest, reinforcing the original pheromone trail as they return.
5. As more and more ants find the food source, the ants eventually lay down a strong pheromone trail to the point that virtually all the ants are directed to the food source.
6. As the food source is depleted, fewer ants are able to find the food, and fewer ants lay down a reinforcing pheromone trail; the pheromone naturally evaporates, and eventually, no ants proceed to the food source, as the ants shift their attention to searching for new food sources.

In an ant colony optimization computational model, the optimization problem is represented as a graph, with nodes representing places and links representing possible paths. An ant colony algorithm mimics ant behavior with simulated ants moving from node to node in the graph, laying down pheromone trails, etc. The process by which ants communicate indirectly by using the environment as an intermediary is known as *stigmergy* (Bonabeau et al. 1999) and is commonly used in agent-based modeling.

Particle swarm optimization (PSO) is another decentralized problem-solving technique in which a swarm of particles is simulated as it moves over a search space in search of a global optimum. A particle stores its best position found so far in

its memory and is aware of the best positions obtained by its neighboring particles. The velocity of each particle adapts over time based on the locations of the best global and local solutions obtained so far, incorporating a degree of stochastic variation in the updating of the particle positions at each iteration.

Artificial Life Algorithms and Agent-Based Modeling

Biologically inspired algorithms are often used with agent-based models. For example, an agent's behavior and its capacity to learn from experience or to adapt to changing conditions can be modeled abstractly through the use of genetic algorithms or neural networks. In the case of a GA, a chromosome effectively represents a single agent action (output) given a specific condition or environmental stimulus (input). Behaviors that are acted on and enable the agent to respond better to environmental challenges are reinforced and acquire a greater share of the chromosome pool. Behaviors that fail to improve the organism's fitness diminish in their representation in the population.

Evolutionary programming can be used to directly evolve programs that represent agent behaviors. For example, Manson (2006) develops a bounded rationality model using evolutionary programming to solve an agent multi-criteria optimization problem.

Artificial neural networks have also been applied to modeling adaptive agent behaviors, in which an agent derives a statistical relationship between the environmental conditions it faces, its history, and its actions, based on feedback on the success or failures of its actions and the actions of others. For example, an agent may need to develop a strategy for bidding in a market, based on the success of its own and other's previous bids and outcomes.

Finally, swarm intelligence approaches are agent based in their basic structure, as described above. They can also be used for system optimization through the selection of appropriate parameters for agent behaviors.

ALife Summary

Based on the previous discussion, the essential features of an ALife program can be summarized as follows:

- *Population*: A population of organisms or individuals is considered. The population may be diversified, and individuals may vary their characteristics, behaviors, and accumulated resources, in both time and space.
- *Interaction*: Interaction requires sensing of the immediate locale, or neighborhood, on the part of an individual. An organism can simply become “aware” of other organisms in its vicinity, or it may have a richer set of interactions with them. The individual also interacts with its (non-agent) environment in its immediate locale. This requirement introduces spatial aspects into the problem, as organisms must negotiate the search for resources through time and space.
- *Sustainment and renewal*: Sustainment and renewal require the acquisition of resources. An organism needs to sense, find, ingest, and metabolize resources or nourishment as an energy source for processing into other forms of nutrients. Resources may be provided by the environment, i.e., outside of the agents themselves, or by other agents. The need for sustainment leads to competition for resources among organisms. Competition could also be a precursor to cooperation and more complex emergent social structures if this proves to be a more effective strategy for survival.
- *Self-reproduction and replacement*: Organisms reproduce by following instructions at least partially embedded within themselves and interacting with the environment and other agents. Passing on traits to the next generation implies a requirement for trait transmission. Trait transmission requires encoding an organism’s traits in a reduced form, that is, a form that contains less than the total information representing the entire organism. It also requires a process for transforming the organism’s traits into a viable set of possible new traits for a new organism. Mutation and cross-over operators enter into such a process. Organisms also leave the population and are replaced by other organisms, possibly with different traits. The organisms can be transformed through changes in their attributes and behaviors, as in, for example, learning or

aging. The populations of organisms can be transformed through the introduction of new organisms and replacement, as in evolutionary adaptation.

As we will see in the section that follows, many of the essential aspects of ALife have been incorporated into the development of agent-based models.

ALife in Agent-Based Modeling

This section briefly touches on the ways in which ALife has motivated agent-based modeling. The form of agent-based models, in terms of their structure and appearance, is directly based on early models from the field of ALife. Several application disciplines in agent-based modeling have been spawned and infused by ALife concepts. Two are covered here. These are how agent-based modeling is applied to social and biological systems.

Agent-Based Modeling Topologies

Agent-based modeling owes much to artificial life in both form and substance. Modeling a population of heterogeneous agents with a diverse set of characteristics is a hallmark of agent-based modeling. The agent perspective is unique among simulation approaches, unlike the process perspective or the state-variable approach taken by other simulation approaches.

As we have seen, agents interact with a small set of neighbor agents in a local area. Agent neighborhoods are defined by how agents are connected, the agent interaction topology. Cellular automata represent agent neighborhoods by using a grid in which the agents exist in the cells, one agent per cell, or as the nodes of the lattice of the grid. The cells immediately surrounding an agent comprise the agent’s neighborhood, and the agents that reside in the neighborhood cells comprise the neighbors. Many agent-based models have been based on this cellular automata spatial representation. The transition from a cellular automaton, such as the Game of Life, to an agent-based model is

accomplished by allowing agents to be distinct from the cells on which they reside and allowing the agents to move from cell to cell across the grid. Agents move according to the dictates of their behaviors, interacting with other agents that happen to be in their local neighborhoods along the way.

Agent interaction topologies have been extended beyond cellular automata to include networks, either predefined and static, as in the case of autocatalytic chemical networks, or endogenous and dynamic, according to the results of agent interactions that occur in the model. Networks allow an agent's neighborhood to be defined more generally and flexibly and, in the case of social agents, more accurately describe social agents' interaction patterns. In addition to cellular automata grids and networks, agent interaction topologies have also been extended across a variety of domains. In summary, agent interaction topologies include:

- Cellular automata grids (agents are cells or are within cells) or lattices (agents are grid points)
- Networks, in which agents are vertices and agent relationships are edges
- Continuous space, in one, two, or three dimensions
- Aspatial random interactions, in which pairs of agents are randomly selected
- Geographical information systems (GIS), in which agents move over geographically defined patches, relaxing the one agent per cell restriction.

Social Agent-Based Modeling

Early social agent-based models were based on ALife's cellular automata approach. In applications of agent-based modeling to social processes, agents represent people or groups of people, and agent relationships represent processes of social interaction (Gilbert and Troitzsch 1999).


Social Agents

Sakoda (1971) formulated one of the first social agent-based models, the checkerboard model, which had some of the key features of a cellular automaton. Following a similar approach,


Schelling developed a model of housing segregation in which agents represent homeowners and neighbors, and agent interactions represent agents' perceptions of their neighbors (Schelling 1971). Schelling studied housing segregation patterns and posed the question of whether it is possible to get highly segregated settlement patterns even if most individuals are, in fact, "color blind." The Schelling model demonstrated that segregated housing areas can develop spontaneously in the sense that system-level patterns can emerge that are not necessarily implied by or consistent with the objectives of the individual agents (Fig. 6). In the model, agents operated according to a fixed set of rules and were not adaptive.

Identifying the social interaction mechanisms for how cooperative behavior emerges among individuals and groups has been addressed using agent-based modeling and evolutionary game theory. Evolutionary game theory accounts for how the repeated interactions of players in a game-theoretic framework affect the development and evolution of the players' strategies. Axelrod showed, using a cellular automata approach, in which agents on the grid employed a variety of different strategies, that a simple tit-for-tat strategy of reciprocal behavior toward individuals is enough to establish sustainable cooperative behavior (Axelrod 1984, 1997). In addition, Axelrod investigated strategies that were self-sustaining and robust in that they reduced the possibility of invasion by agents having other strategies.

Epstein and Axtell introduced the notion of an external environment that agents interact with in addition to other agents. In their groundbreaking Sugarscape model of artificial societies, agents interacted with their environment depending on their location in the grid (Epstein and Axtell 1996). This allowed agents to access environmental variables, extract resources, etc., based on the location. In numerous computational experiments, Sugarscape agents emerged with a variety of characteristics and behaviors, highly suggestive of a realistic, although rudimentary and abstract, society (Fig. 7). They observed emergent processes that they interpreted as death, disease, trade, wealth, sex and reproduction, culture,



Agent-Based Modeling and Artificial Life, Fig. 6 Schelling housing segregation model



Agent-Based Modeling and Artificial Life, Fig. 7 Sugarscape artificial society simulation in the Repast agent-based modeling toolkit

conflict, and war, as well as externalities such as pollution. As agents interacted with their neighbors as they moved around the grid, the interactions resulted in a contact network, that is, a network consisting of nodes and links. The nodes are agents, and the links indicate the agents that have been neighbors at some point in the course of their movements over the grid. Contact networks were the basis for studying contagion and epidemics in the Sugarscape model. Understanding the agent rules that govern how networks are structured and grow, how quickly information is communicated through networks, and the kinds of relationships that networks embody is an important aspect of modeling agents.

Culture and Generative Social Science

Dawkins, who has written extensively on aspects of Darwinian evolution, coined the term *meme* as the smallest element of culture that is transmissible between individuals, similar to the notion of the gene as being the primary unit of transmitting genetic information (Dawkins 1989). Several social agent-based models are based on a meme representation of culture as shared or collective agent attributes.

In the broadest terms, social agent-based simulation is concerned with social interaction and social processes. Emergence enters into social simulation through *generative social science* whose goal is to model social processes as emergent processes and their emergence as the result of social interactions. Epstein has argued that social processes are not fully understood unless one is able to theorize how they work at a deep level and have social processes emerge as part of a computational model (Epstein 2007). More recent work has treated culture as a fluid and dynamic process subject to interpretation of individual agents, more complex than the genotype/phenotype framework would suggest.

ALife and Biology


ALife research has motivated many agent-based computational models of biological systems, and at all scales, ranging from the cellular level, or even the subcellular molecular level, as the basic unit of agency, to complex organisms embedded in larger structures such as food webs or complex ecosystems.

From Cellular Automata to Cells

Cellular automata are a natural application to modeling cellular systems (Alber et al. 2003; Ermentrout and Edelstein-Keshet 1993). One approach uses the cellular automata grid and cells to model structures of stationary cells comprising a tissue matrix. Each cell is a tissue agent. Mobile cells consisting of pathogens and antibodies are also modeled as agents. Mobile agents diffuse through tissue and interact with tissue and other colocated mobile cells. This approach is the basis for agent-based models of the immune system. Celada and Seiden (1992) used bit strings to model the cell receptors in a cellular automaton model of the immune system called IMMSIM. This approach was extended to a more general agent-based model and implemented to maximize the number of cells that could be modeled in the CIMMSIM and ParlImm systems (Bernaschi and Castiglione 2001). The Basic Immune Simulator uses a general agent-based framework (the Repast agent-based modeling toolkit) to model the interactions between the cells of the innate and adaptive immune system (Folcik et al. 2007). These approaches for modeling the immune system have inspired several agent-based models of intrusion detection for computer networks (see, e.g., Azzedine et al. 2007) and have found use in modeling the development and spread of cancer (Preziosi 2003).

At the more macroscopic level, agent-based epidemic models have been developed using network topologies. These models include people and some representation of pathogens as individual agents for natural (Bobashev et al. 2007) and potentially man-made (Carley et al. 2006) epidemics.

Modeling bacteria and their associated behaviors in their natural environments are another direction of agent-based modeling. Expanding beyond the basic cellular automata structure into continuous space and network topologies, Emonet et al. (2005) developed AgentCell, a multi-scale agent-based model of *E. coli* bacteria motility (Fig. 8). In this multi-scale agent-based simulation, molecules within a cell are modeled as individual agents. The molecular reactions comprising the signal transduction network for chemotaxis are modeled using an embedded stochastic simulator, StochSim (Le Novère and



Agent-Based Modeling and Artificial Life, Fig. 8 AgentCell multi-scale agent-based model of bacterial chemotaxis

Shimizu 2001). This multi-scale approach allows the motile (macroscopic) behavior of colonies of bacteria to be modeled as a direct result of the modeled microlevel processes of protein production within the cells, which are based on individual molecular interactions.

Artificial Ecologies

Early models of ecosystems used approaches adapted from physical modeling, especially models of idealized gases based on statistical mechanics. More recently, individual-based models have been developed to represent the full range of individual diversity by explicitly modeling individual attributes or behaviors and aggregating across individuals for an entire population (DeAngelis and Gross 1992). Agent-based approaches model a diverse set of agents and their interactions based on their relationships, incorporating adaptive behaviors as appropriate. For example, food webs represent the complex, hierarchical network of agent relationships in local ecosystems (Peacor et al. 2006). Agents are individuals or species representatives. Adaptation and learning for agents in such food webs can be modeled to explore diversity, relative population sizes, and resiliency to environmental insult.

Adaptation and Learning in Agent-Based Models Biologists consider adaptation to be an essential part of the process of evolutionary change. Adaptation occurs at two levels: the individual level and the population level. In parallel with these notions, agents in an ABM adapt by changing their individual behaviors or by changing their proportional representation in the population. Agents adapt their behaviors at the individual level through learning from experience in their modeled environment.

With respect to agent-based modeling, theories of learning by individual agents or collectives of agents, as well as algorithms for how to model learning, become important. Machine learning is a field consisting of algorithms for recognizing patterns in data (such as data mining) through techniques such as supervised learning, unsupervised learning, and reinforcement learning (Alpaydm 2004; Bishop 2007). Genetic algorithms (Goldberg 1989) and related techniques such as learning classifier systems (Holland et al. 2000) are commonly used to represent agent learning in agent-based models. In ABM applications, agents learn through interactions with the simulated environment in which they are embedded as the simulation precedes through time, and agents modify their behaviors accordingly.

Agents may also adapt collectively at the population level. Those agents having behavioral rules better suited to their environments survive and thrive, and those agents not so well suited are gradually eliminated from the population.

Future Directions

Agent-based modeling continues to be inspired by ALife – in the fundamental questions it is trying to answer, in the algorithms that it employs to model agent behaviors and solve agent-based models, and in the computational architectures that are employed to implement agent-based models. The future of the fields of both ALife and ABM will continue to be intertwined in essential ways in the coming years.

Computational advances will continue at an ever-increasing pace, opening new vistas for computational possibilities in terms of expanding the scale of models that are possible. Computational advances will take several forms, including advances in computer hardware including new chip designs, multi-core processors, and advanced integrated hardware architectures. Software that take advantage of these designs and in particular computational algorithms and modeling techniques and approaches will continue to provide opportunities for advancing the scale of applications and allow more features to be included in agent-based models as well as ALife applications. These will be opportunities for advancing applications of ABM to ALife in both the realms of scientific research and in policy analysis.

Real-world optimization problems routinely solved by business and industry will continue to be solved by ALife-inspired algorithms. The use of ALife-inspired agent-based algorithms for solving optimization problems will become more widespread because of their natural implementation and ability to handle ill-defined problems.

Emergence is a key theme of ALife. ABM offers the capability to model the emergence of order in a variety of complex and complex adaptive systems. Inspired by ALife, identifying the fundamental mechanisms responsible for higher-order emergence and exploring these with agent-based modeling will be an important and promising research area.

Advancing social sciences beyond the genotype/phenotype framework to address the generative nature of social systems in their full complexity is a requirement for advancing computational social models. Recent work has treated culture as a fluid and dynamic process subject to interpretation of individual agents, more complex in many ways than that provided by the genotype/phenotype framework.

Agent-based modeling will continue to be the avenue for exploring new constructs in ALife. If true artificial life is ever developed in silico, it will most likely be done using the methods and tools of agent-based modeling.

Bibliography

Primary Literature

- Adami C (1998) Introduction to artificial life. TELOS, Santa Clara
- Alber MS, Kiskowski MA, Glazier JA, Jiang Y (2003) On cellular automaton approaches to modeling biological cells. In: Rosenthal J, Gilliam DS (eds) Mathematical systems theory in biology, communication, and finance, IMA volume. Springer, New York, pp 1–39
- Alpaydm E (2004) Introduction to machine learning. MIT Press, Cambridge
- Axelrod R (1984) The evolution of cooperation. Basic Books, New York
- Axelrod R (1997) The complexity of cooperation: agent-based models of competition and collaboration. Princeton University Press, Princeton
- Azzedine B, Renato BM, Kathia RLJ, Joao Bosco MS, Mirela SMAN (2007) An agent based and biological inspired real-time intrusion detection and security model for computer network operations. *Comput Commun* 30(13):2649–2660
- Back T (1996) Evolutionary algorithms in theory and practice: evolution strategies, evolutionary programming, genetic algorithms. Oxford University Press, New York
- Berlekamp ER, Conway JH, Guy RK (2003) Winning ways for your mathematical plays, 2nd edn. AK Peters, Natick
- Bernaschi M, Castiglione F (2001) Design and implementation of an immune system simulator. *Comput Biol Med* 31(5):303–331
- Bishop CM (2007) Pattern recognition and machine learning. Springer, New York
- Bobashev GV, Goedecke DM, Yu F, Epstein JM (2007) A hybrid epidemic model: combining the advantages of agent-based and equation-based approaches. In: Henderson SG, Biller B, Hsieh M-H, Shortle J, Tew JD, Barton RR (eds) *Proceeding 2007 winter simulation conference*, Washington, pp 1532–1537

- Bonabeau E (1997) From classical models of morphogenesis to agent-based models of pattern formation. *Artif Life* 3:191–211
- Bonabeau E, Dorigo M, Theraulaz G (1999) *Swarm intelligence: from natural to artificial systems*. Oxford University Press, New York
- Carley KM, Fridsma DB, Casman E, Yahja A, Altman N, Chen LC, Kaminsky B, Nave D (2006) Biowar: scalable agent-based model of bioattacks. *IEEE Trans Syst Man Cybern Part A: Syst Hum* 36(2):252–265
- Celada F, Seiden PE (1992) A computer model of cellular interactions in the immune system. *Immunol Today* 13(2):56–62
- Clerc M (2006) *Particle swarm optimization*. ISTE Publishing, London
- Dawkins R (1989) *The selfish gene*, 2nd edn. Oxford University Press, Oxford
- DeAngelis DL, Gross LJ (eds) (1992) Individual-based models and approaches in ecology: populations, communities and ecosystems. *Proceedings of a symposium/workshop*, Knoxville, 16–19 May 1990. Chapman & Hall, New York. ISBN 0-412-03171-X
- Dorigo M, Stützle T (2004) *Ant colony optimization*. MIT Press, Cambridge
- Dreyfus HL (1979) *What computers can't do: the limits of artificial intelligence*. Harper & Row, New York
- Eiben AE, Smith JE (2007) *Introduction to evolutionary computing*, 2nd edn. Springer, New York
- Eigen M, Schuster P (1979) *The hypercycle: a principle of natural self-organization*. Springer, Berlin
- Emonet T, Macal CM, North MJ, Wickersham CE, Cluzel P (2005) AgentCell: a digital single-cell assay for bacterial chemotaxis. *Bioinformatics* 21(11):2714–2721
- Engelbrecht AP (2006) *Fundamentals of computational swarm intelligence*. Wiley, Hoboken
- Epstein JM (2007) *Generative social science: studies in agent-based computational modeling*. Princeton University Press, Princeton
- Epstein JM, Axtell R (1996) *Growing artificial societies: social science from the bottom up*. MIT Press, Cambridge
- Ermentrout GB, Edelstein-Keshet L (1993) Cellular automata approaches to biological modeling. *J Theor Biol* 160(1):97–133
- Fogel LJ, Owens AJ, Walsh MJ (1966) *Artificial intelligence through simulated evolution*. Wiley, Hoboken
- Folcik VA, An GC, Orosz CG (2007) The basic immune simulator: an agent-based model to study the interactions between innate and adaptive immunity. *Theor Biol Med Model* 4(39):1–18. <http://www.tbiomed.com/content/4/1/39>
- Fontana W (1992) Algorithmic chemistry. In: Langton CG, Taylor C, Farmer JD, Rasmussen S (eds) *Artificial life II: proceedings of the workshop on artificial life*, Santa Fe, Feb 1990, Santa Fe Institute studies in the sciences of the complexity, vol X. Addison-Wesley, Reading, pp 159–209
- Gardner M (1970) The fantastic combinations of John Conway's new solitaire game life. *Sci Am* 223:120–123
- Gilbert N (2002) Varieties of emergence. In: Macal C, Sallach D (eds) *Proceedings of the agent 2002 conference on social agents: ecology, exchange and evolution*, Chicago, 11–12 Oct 2002, pp 1–11. Available on CD and at www.agent2007.anl.gov
- Gilbert N, Troitzsch KG (1999) *Simulation for the social scientist*. Open University Press, Buckingham
- Goldberg DE (1989) *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley, Reading
- Goldberg DE (1994) Genetic and evolutionary algorithms come of age. *Commun ACM* 37(3):113–119
- Holland JH (1975) *Adaptation in natural and artificial systems*. University of Michigan, Ann Arbor
- Holland J (1995) *Hidden order: how adaptation builds complexity*. Addison-Wesley, Reading
- Holland JH, Booker LB, Colombetti M, Dorigo M, Goldberg DE, Forrest S, Riolo RL, Smith RE, Lanzi PL, Stolzmann W, Wilson SW (2000) What is a learning classifier system? In: Lanzi PL, Stolzmann W, Wilson SW (eds) *Learning classifier systems, from foundations to applications*. Springer, London, pp 3–32
- Kauffman SA (1993) *The origins of order: self-organization and selection in evolution*. Oxford University Press, Oxford
- Koza JR (1992) *Genetic programming: on the programming of computers by means of natural selection*. MIT Press, Cambridge, 840 pp
- Langton CG (1984) Self-reproduction in cellular automata. *Physica D* 10:135–144
- Langton CG (1986) Studying artificial life with cellular automata. *Physica D* 22:120–149
- Langton CG (1989a) Preface. In: Langton CG (ed) *Artificial life: proceedings of an interdisciplinary workshop on the synthesis and simulation of living systems*, Los Alamos, Sept 1987, Addison-Wesley, Reading, pp xv–xxvi
- Langton CG (1989b) Artificial life. In: Langton CG (ed) *Artificial life: the proceedings of an interdisciplinary workshop on the synthesis and simulation of living systems*, Los Alamos, Sept 1987, Santa Fe Institute studies in the sciences of complexity, vol VI. Addison-Wesley, Reading, pp 1–47
- Langton CG (1992) Life at the edge of chaos. In: Langton CG, Taylor C, Farmer JD, Rasmussen S (eds) *Artificial life II: proceedings of the workshop on artificial life*, Santa Fe, Feb 1990, Santa Fe Institute studies in the sciences of the complexity, vol X. Addison-Wesley, Reading, pp 41–91
- Le Novère N, Shimizu TS (2001) Stochsim: modelling of stochastic biomolecular processes. *Bioinformatics* 17(6):575–576
- Lindenmeyer A (1968) Mathematical models for cellular interaction in development. *J Theor Biol* 18:280–315
- Lucas JR (1961) Minds, machines and godel. *Philosophy* 36(137):112–127
- Manson SM (2006) Bounded rationality in agent-based models: experiments with evolutionary programs. *Int J Geogr Inf Sci* 20(9):991–1012
- Mehrotra K, Mohan CK, Ranka S (1996) *Elements of artificial neural networks*. MIT Press, Cambridge
- Mitchell M, Forrest S (1994) Genetic algorithms and artificial life. *Artif Life* 1(3):267–289
- Ofria C, Wilke CO (2004) Avida: a software platform for research in computational evolutionary biology. *Artif Life* 10(2):191–229

- Olariu S, Zomaya AY (eds) (2006) Handbook of bioinspired algorithms and applications. Chapman, Boca Raton, p 679
- Padgett JF, Lee D, Collier N (2003) Economic production as chemistry. *Ind Corp Chang* 12(4):843–877
- Peacor SD, Riolo RL, Pascual M (2006) Plasticity and species coexistence: modeling food webs as complex adaptive systems. In: Pascual M, Dunne JA (eds) *Ecological networks: linking structure to dynamics in food webs*. Oxford University Press, New York, pp 245–270
- Penrose R (1989) The emperor's new mind: concerning computers, minds, and the laws of physics. Oxford University Press, Oxford
- Poundstone W (1985) The recursive universe. Contemporary Books, Chicago, 252 pp
- Preziosi L (ed) (2003) Cancer modelling and simulation. Chapman, Boca Raton
- Ray TS (1991) An approach to the synthesis of life (terra simulator). In: Langton CG, Taylor C, Farmer JD, Ras-mussen S (eds) *Artificial life II: proceedings of the workshop on artificial life*. Wesley, Redwood City, pp 371–408
- Rechenberg I (1973) Evolutionsstrategie: optimierung Technischer Systeme Nach Prinzipien Der Biologischen evolution. Frommann-Holzboog, Stuttgart
- Sakoda JM (1971) The checkerboard model of social interaction. *J Math Soc* 1:119–132
- Schelling TC (1971) Dynamic models of segregation. *J Math Soc* 1:143–186
- Searle JR (1990) Is the brain a digital computer? Presidential Address to the American Philosophical Association
- Taub AH (ed) (1961) John Von Neumann: collected works. vol V: Design of computers, theory of automata and numerical analysis (Delivered at the Hixon Symposium, Pasadena, Sept 1948). Pergamon Press, Oxford
- Turing AM (1938) On computable numbers with an application to the entscheidungsproblem. *Process Lond Math Soc* 2(42):230–265
- Turing AM (1952) The chemical basis of morphogenesis. *Philos Trans Royal Soc B* 237:37–72
- von Neumann J (1966) In: Burks AW (ed) *Theory of self-reproducing automata*. University of Illinois Press, Champaign
- Wilke CO, Adami C (2002) The biology of digital organisms. *Trends Ecol Evol* 17(11):528–532
- Wilke CO, Chow SS (2006) Exploring the evolution of ecosystems with digital organisms. In: Pascual M, Dunne JA (eds) *Ecological networks: linking structure to dynamics in food webs*. Oxford University Press, New York, pp 271–286
- Wolfram S (1984) Universality and complexity in cellular automata. *Physica D* 1–35
- Wolfram S (1999) *The mathematica book*, 4th edn. Wolfram Media/Cambridge University Press, Champaign
- Wolfram S (2002) *A new kind of science*. Wolfram Media, Champaign
- Batty M (2007) *Cities and complexity: understanding cities with cellular automata, agent-based models, and fractals*. MIT Press, Cambridge
- Bedau MA (2002) The scientific and philosophical scope of artificial life. *Leonardo* 35:395–400
- Bedau MA (2003) *Artificial life: organization, adaptation and complexity from the bottom up*. *TRENDS Cognit Sci* 7(11):505–512
- Copeland BJ (2004) *The essential turing*. Oxford University Press, Oxford, 613 pp
- Ganguly N, Sikdar BK, Deutsch A, Canright G, Chaudhuri PP (2008) A survey on cellular automata. www.cs.unibo.it/bison/publications/CAsurvey.pdf
- Griffeth D, Moore C (eds) (2003) *New constructions in cellular automata*, Santa Fe Institute studies in the sciences of complexity proceedings. Oxford University Press, New York, 360 pp
- Gutowitz H (ed) (1991) *Cellular automata: theory and experiment*. Special issue of *Physica D*. 499 pp
- Hraber T, Jones PT, Forrest S (1997) The ecology of echo. *Artif Life* 3:165–190
- International Society for Artificial Life web page (2008) www.alife.org. Accessed 8 Mar 2008
- Jacob C (2001) *Illustrating evolutionary computation with mathematica*. Academic, San Diego, 578 pp
- Michael CF, Fred WG, Jay A (2005) Simulation optimization: a review, new developments, and applications. In: *Proceedings of the 37th conference on Winter simulation*, Orlando
- Miller JH, Page SE (2007) *Complex adaptive systems: an introduction to computational models of social life*. Princeton University Press, Princeton
- North MJ, Macal CM (2007) *Managing business complexity: discovering strategic solutions with agent-based modeling and simulation*. Oxford University Press, New York
- Pascual M, Dunne JA (eds) (2006) *Ecological networks: linking structure to dynamics in food webs*, Santa Fe Institute studies on the sciences of complexity. Oxford University Press, New York
- Simon H (2001) *The sciences of the artificial*. MIT Press, Cambridge
- Sims K (1991) Artificial evolution for computer graphics. *ACM SIGGRAPH '91*, Las Vegas, July 1991, pp 319–328
- Sims K (1994) Evolving 3D morphology and behavior by competition. *Artif Life* IV:28–39
- Terzopoulos D (1999) Artificial life for computer graphics. *Commun ACM* 42(8):33–42
- Toffoli T, Margolus N (1987) *Cellular automata machines: a new environment for modeling*. MIT Press, Cambridge, 200 pp
- Tu X, Terzopoulos D (1994) Artificial fishes: physics, locomotion, perception, behavior. In: *Proceedings of SIGGRAPH '94*, 24–29 July 1994, Orlando, pp 43–50
- Weisbuch G (1991) *Complex systems dynamics: an introduction to automata networks*, translated from French by Ryckebusch S. Addison-Wesley, Redwood City
- Wiener N (1948) *Cybernetics, or control and communication in the animal and the machine*. Wiley, New York
- Wooldridge M (2000) *Reasoning about rational agents*. MIT Press, Cambridge

Books and Reviews

- Artificial Life (journal) web page (2008) <http://www.mitpressjournals.org/loi/artl>. Accessed 8 Mar 2008
- Banks ER (1971) *Information processing and transmission in cellular automata*. PhD dissertation, Massachusetts Institute of Technology

Embodied and Situated Agents, Adaptive Behavior in

Stefano Nolfi

Institute of Cognitive Sciences and Technologies, National Research Council (CNR), Rome, Italy

Article Outline

Definition of the Subject
Introduction
Embodiment and Situatedness
Behavior and Cognition as Complex Adaptive Systems
Behavior and Cognition as Phenomena
 Originating from the Interaction Between Coupled Dynamical Processes
Behavior and Cognition as Phenomena with a Multilevel and Multi-scale Organization
On the Top-Down Effect from Higher to Lower Levels of Organization
Adaptive Methods
Evolutionary Robotics Methods
Developmental Robotics Methods
The Incremental Nature of the Developmental Process
The Social Nature of the Developmental Process
Exploitation of the Interaction Between Concurrent Developmental Processes
Discussion and Conclusion
Bibliography

Glossary

Embodied agent Indicates an artificial system (simulated or physical) which has body (characterized by physical properties such as shape, dimension, weight, etc.), actuators (e.g., motorized wheels, motorized articulated joints), and sensors (e.g., touch sensors or

vision sensors). For a more restricted definition, see the concluding section of the entry.

Morphological computation Indicates the ability of the body of an agent (with certain specific characteristics) to control its interaction with the environment so to produce a given desired behavior.

Ontogenesis Indicates the variations which occur in the phenotypical characteristics of an artificial agent (i.e., in the characteristics of the control system or of the body of the agent), while it interacts with the environment.

Phylogenesis Indicates the variations of the genetic characteristics of a population of artificial agents throughout generations.

Situated agent Indicates an artificial system which is located in a physical environment (simulated or real) with which it interacts on the basis of the law of physics. For a more restricted definition, see the concluding section of the entry.

Definition of the Subject

Adaptive behavior concerns the study of how organisms develop their behavioral and cognitive skills through a synthetic methodology which consists in designing artificial agents which are able to adapt to their environment autonomously. These studies are important both from a modeling point of view (i.e., for making progress in our understanding of intelligence and adaptation in natural beings) and from an engineering point of view (i.e., for making progresses in our ability to develop artifacts displaying effective behavioral and cognitive skills).

Introduction

Adaptive behavior research concerns the study of how organisms can develop behavioral and cognitive skills by adapting to the environment and

to the task they have to fulfill autonomously (i.e., without human intervention). This goal is achieved through a synthetic methodology, i.e., through the synthesis of artificial creatures which (i) have a body, (ii) are situated in an environment with which they interact, and (iii) have characteristics which vary during an adaptation process. In the rest of the entry, we will use the term “agent” to indicate artificial creatures which possess the first two features described above and the term “adaptive agent” to indicate artificial creatures which also possess the third feature.

The agents and the environment might be simulated or real. In the former case, the characteristics of agents’ body, motor, and sensory system, the characteristics of the environment, and the rules that regulate the interactions between all the elements are simulated on a computer. In the latter case, the agents consist of physical entities (mobile robots) situated in a physical environment with which they interact on the basis of the physical laws.

The adaptive process which regulates how the characteristics of the agents (and eventually of the social environment) change might consist of a population-based evolutionary process and/or of a developmental/learning process. In the former case, the characteristics of the agents do not vary during their “lifetime” (i.e., during the time in which the agents interact with the environment) but, phylogenetically, while individual agents “reproduce.” In the latter case, the characteristics of the agents vary ontogenetically, while they interact with the environment. The criteria which determine how variations are generated and/or whether or not variations are retained can be task dependent and/or task independent, i.e., might be based on an evaluation of whether the variation increases or decreases agents’ ability to display a behavior which is adapted to the task/environment or might be based on task-independent criteria (i.e., general criteria which do not reward directly the exhibition of the requested skill).

The entry is organized as follows. In section “[Embodiment and Situatedness](#),” we briefly introduce the notion of *embodiment* and *situatedness* and their implications. In section

“[Behavior and Cognition as Complex Adaptive Systems](#),” we claim that behavior and cognition in embodied and situated adaptive agents should be characterized as a complex adaptive system. In section “[Adaptive Methods](#),” we briefly describe the methods which can be used to synthesize embodied and situated adaptive agents. Finally in section “[Discussion and Conclusions](#),” we draw our conclusions.

Embodiment and Situatedness

The notion of *embodiment* and *situatedness* has been introduced (Brooks 1991a, b; Clark 1997; Pfeifer and Bongard 2007; Varela et al. 1991) to characterize systems (e.g., natural organism and robots) which have a physical body and which are situated in a physical environment with which they interact. In this and in the following sections, we will briefly discuss the general implications of these two fundamental properties. This analysis will be further extended in the concluding section where we will claim on the necessity to distinguish between a weak and a strong notion of embodiment and situatedness.

One first important implication of being embodied and situated consists in the fact that these agents and their parts are characterized by their physical properties (e.g., weight, dimension, shape, elasticity, etc.), are subjected to the laws of physics (e.g., inertia, friction, gravity, energy consumption, deterioration, etc.), and interact with the environment through the exchange of energy and physical material (e.g., forces, sound waves, light waves, etc.). Their physical nature also implies that they are quantitative in state and time (van Gelder 1998). The fact that these agents are quantitative in time implies, for example, that the joints which connect the parts of a robotic arm can assume any possible position within a given range. The fact that these agents are quantitative in time implies, for example, that the effects of the application of a force to a joint depend from the time duration of its application.

One second important implication is that the information measured by the sensors is not only a function of the environment but also of the

relative position of the agent in the environment. This implies that the motor actions performed by an agent, by modifying the agent/environmental relation or the environment, co-determine the agent sensory experiences.

One third important implication is that the information extracted by the sensors from the external environment is egocentric (depends from the current position and the orientation of the agent in the environment), local (only provide information related to the local observable portion of the environment), incomplete (e.g., due to visual occlusion), and subjected to noise. Similar characteristics apply to the motor actions produced by the agent's effectors.

It is important to notice that these characteristics do not only represent constraints but also opportunities to be exploited. Indeed, as we will see in the next section, the exploitation of some of these characteristics might allow embodied and situated agents to solve their adaptive problems through solutions which are robust and parsimonious (i.e., minimal) with respect to the complexity of the agent's body and control system.

Behavior and Cognition as Complex Adaptive Systems

In embodied and situated agents, behavioral and cognitive skills are dynamical properties which unfold in time and which arise from the interaction between agents' nervous system, body, and the environment (Beer 1995; Chiel and Beer 1997; Keijzer 2001; Nolfi 2005; Nolfi and Floreano 2000) and from the interaction between dynamical processes occurring within the agents' control system, the agents' body, and within the environment (Beer 2003; Gigliotta and Nolfi 2008; Tani and Fukumura 1997). Moreover, behavioral and cognitive skills typically display a multilevel and multi-scale organization involving bottom-up and top-down influences between entities at different levels of organization. These properties imply that behavioral and cognitive skills in embodied and situated agents can be properly characterized as complex adaptive systems (Nolfi 2005).

These aspects and the complex system nature of behavior and cognition will be illustrated in more details in the next subsections also with the help of examples. The theoretical and practical implication of these aspects for developing artificial agents able to exhibit effective behavioral and cognitive skills will be discussed in the forthcoming sections.


Behavior and Cognition as Emergent Dynamical Properties

Behavior and cognition are dynamical properties which unfold in time and which emerge from high-frequent nonlinear interactions between the agent, its body, and the external environment (Chiel and Beer 1997).

At any time step, the environmental and the agent/environmental relation co-determine the body and the motor reaction of the agent which, in turn, co-determines how the environment and/or the agent/environmental relation vary. Sequences of these interactions, occurring at a fast-time rate, lead to a dynamical process – behavior – which extends over significant larger time span than the interactions (Fig. 1).

Since interactions between the agent's control system, the agent's body, and the external environment are nonlinear (i.e., small variations in sensory states might lead to significantly different motor actions) and dynamical (i.e., small variations in the action performed at time t might significantly impact later interactions at time (t_{+x})), the relation between the rules that govern the interactions and the behavioral and cognitive skills originating from the interactions tend to be very indirect. Behavioral and cognitive skills thus emerge from the interactions between the three foundational elements and cannot be traced back to any of the three elements taken in isolation. Indeed, the behavior displayed by an embodied and situated agent can hardly be predicted or inferred from an external observer even on the basis of a complete knowledge of the interacting elements and of the rules governing the interactions.

A clear example of how behavioral skill might emerge from the interaction between the agents'



Embodied and Situated Agents, Adaptive Behavior in, Fig. 1 A schematic representation of the relation between agent's control system, agent's body, and the environment. The behavioral and cognitive skills displayed by the agent are the emergent result of the bidirectional interactions (represented with *full arrows*) between the three constituting elements – agent's control system, agent's body, and environment. The *dotted arrows* indicate that the three constituting elements

might be dynamical systems on their own. In this case, agents' behavioral and cognitive skills not only result from the dynamics originating from the agent/body/environmental interactions but also from the combination and the interaction between dynamical processes occurring within the agent's body, within the agent's control system, and within the environment (see section “[Behavior and Cognition as Complex Adaptive Systems](#)”)


body and the environment is constituted by the passive walking machines developed in simulation by McGeer (1990) – a two-dimensional bipedal machines able to walk down a 4° slope with no motors and no control system (Fig. 2). The walking behavior arises from the fact that the physical forces resulting from gravity and from the collision between the machine and the slope produce a movement of the robot and the fact that robot's movements produce a variation of the agent-environmental relation which in turn produce a modification of the physical forces to which the machine will be subjected in the next time step. The sequence of bidirectional effects between the robot's body and the environment can lead to a stable dynamical process – the walking behavior.

The type of behavior which arises from the robot/environmental interaction depends from

the characteristics of the environment, the physics law which regulate the interaction between the body and the environment, and the characteristics of the body. The first two factors can be considered as fixed, but the third factor, the body structure, can be adapted to achieve a given function. Indeed, in the case of this biped robot, the author carefully selected the leg length, the leg mass, and the foot size to obtain the desired walking behavior. In more general term, this example shows how the role of regulating the interaction between the robot and the environment in the appropriate way can be played not only from the control system but also from the body itself providing that the characteristics of the body have been shaped so to favor the exhibition of the desired behavior. This property, i.e., the ability of the body to control its interaction with the environment, has been named with the

Embodied and Situated Agents, Adaptive Behavior in, Fig. 2

A schematization of the passive walking machine developed by McGeer (1990). The machine includes two passive knee joints and a passive hip joint



term “morphological computation” (Pfeifer et al. 2006). For related work which demonstrates how effective walking machines can be obtained by integrating passive walking techniques with simple control mechanisms, see Bongard and Paul (2001), Endo et al. (2002), and Vaughan et al. (2004). For related works which show the role of elastic material and elastic actuators for morphological computing, see Massera et al. (2007) and Schmitz et al. (2007).


To illustrate how behavioral and cognitive skills might emerge from agent’s body, agent’s control system, and environmental interactions, we describe a simple experiment in which a small-wheeled robot situated in an arena surrounded by walls has been evolved to find and to remain close to a cylindrical object. The Khepera robot (Mondada et al. 1993) is provided with eight infrared sensors and two motors controlling the two corresponding wheels (Fig. 3).

From the point of view of an external observer, solving this problem requires robots able to: (a) explore the environment until an obstacle is detected, (b) discriminate whether the obstacle detected is a wall or a cylindrical object, and (c) approach or avoid objects depending on the object type. Some of these behaviors (e.g., the wall-avoidance behavior) can be obtained through simple control mechanisms, but others require nontrivial control mechanisms. Indeed, a detailed analysis of the sensory patterns experienced by the robot indicated that the task of

discriminating the two objects is far from trivial since the two classes of sensory patterns experienced by robots close to a wall and close to cylindrical objects largely overlap.

The attempt to solve this problem through an evolutionary adaptive method (see section “Adaptive Methods”) in which the free parameters (i.e., the parameters which regulate the fine-grained interaction between the robot and the environment) are varied randomly and in which variations are retained or discarded on the basis of an evaluation of the overall ability of the robot (i.e., on the basis of the time spent by the robot close to the cylindrical object) demonstrated how adaptive robots can find solutions which are robust and parsimonious in terms of control mechanisms (Nolfi 2002). Indeed, in all replications of this experiment, evolved robot solves the problem by moving forward, by avoiding walls, and by oscillating back and forth and left and right close to cylindrical objects (Fig. 3, right). All these behaviors result from sequences of interactions between the robot and the environment mediated by four types of simple control rules which consist in: turning left when the right infrared sensors are activated, turning right when the left infrared sensors are activated, moving back when the frontal infrared sensors are activated, and moving forward when the frontal infrared sensors are not activated.

To understand how these simple control rules can produce the required behaviors and the



Embodied and Situated Agents, Adaptive Behavior in, Fig. 3 *Left*, the agent situated in the environment. The agent is a Khepera robot (Mondada et al. 1993). The environment consists of an arena of 60×35 cm containing cylindrical objects placed in a randomly selected location. *Right*, angular trajectories of an evolved robot close to a wall (*top graph*) and to a cylinder (*bottom graph*). The picture was obtained by placing the robot at a


random position in the environment, leaving it free to move for 500 time steps each lasting 100 ms, and recording its relative movements with respect to the two types of objects for distances smaller than 45 mm. The x -axis and the y -axis indicate the relative angle (in degrees) and distance (in mm) between the robot and the corresponding object. For the sake of clarity, *arrows* are used to indicate the relative direction but not the amplitude of movements

required arbitration between behaviors, we should consider that the same motor responses produce different effects on different agent/environmental situations. For example, the execution of a left-turning action close to a cylindrical object and the subsequent modification of the robot/object relative position produce a new sensory state which triggers a right-turning action. Then, the execution of the latter action and the subsequent modification of the robot/object relative position produce a new sensory state which triggers a left-turning action. The combination and the alternation of these left- and right-turning actions over time produce an attractor in the agent/environmental dynamics (Fig. 3, right, bottom graph) which allows the robot to remain close to the cylindrical object. On the other hand, the execution of a left-turning behavior close to a wall object and the subsequent modification of the robot/wall position produce a new sensory state which triggers the reiteration of the same motor action. The execution of a sequence of left-turning action then leads to the avoidance of the object and to a modification of the robot/environmental relation which finally lead to a perception of a sensory state which trigger a move-forward behavior (Fig. 4, right, top graph).

Before concluding the description of this experiment, it is important to notice that although

the rough classification of the robot motor responses into four different types of actions is useful to describe the strategy with which these robots solve the problem qualitatively, the quantitative aspects which characterize the robot motor reactions (e.g., how sharply a robot turns given a certain pattern of activation of the infrared sensors) are crucial for determining whether the robot will be able to solve the problem or not. Indeed, small differences in the robot's motor response tend to cumulate in time and might prevent the robot for producing successful behavior (e.g., might prevent the robot to produce a behavioral attractor close to cylindrical objects).

This experiment clearly exemplifies some important aspects which characterize all adaptive behavioral system, i.e., systems which are embodied and situated and which have been designed or adapted so to exploit the properties that emerge from the interaction between their control system, their body, and the external environment. In particular, it demonstrates how required behavioral and cognitive skills (i.e., object categorization skills) might emerge from the fine-grained interaction between the robot's control system, body, and the external environment without the need of dedicated control mechanisms. Moreover, it demonstrates how the relation between the control rules which mediate



Embodied and Situated Agents, Adaptive Behavior in, Fig. 5 The state of the two internal neurons (i_1 and i_2) of the robot recorded for 330 s while the robot performs about five laps of the environment. The s , a , b , c , and d labels indicate the internal states corresponding to five different positions of the robot in the environment shown in Fig. 4. The other labels indicate the position of the fixed

point attractors in the robot's internal dynamics corresponding to five types of sensory states experienced by the robot when it detects a light in its frontal side (LF), a light on its rear side (LR), an obstacle on its right and frontal side (OFR), an obstacle on its right side (OR), and no obstacles and no lights (NO)

should be turned on when the robot revisits the environmental area in which the black disk was previously found. The robot's controller consists of a three-layer neural network which includes a layer of sensory neurons (which encode the state of the corresponding sensors), a layer of motor neurons which encode the state of the actuators, and a layer of internal neurons which consist of leaky integrators operating at tunable timescale (Beer 1995; Gigliotta and Nolfi 2008). The free parameters of the robot's neural controllers (i.e., the connection weights and the time constant of the internal neurons which regulate the time rate at which these neurons change their state over-time) were adapted through an evolutionary technique (Nolfi and Floreano 2000).

By analyzing the evolved robot, the authors observed how they are able to generate a spatial representation of the environment and of their location in the environment while they are situated in the environment itself. Indeed, while the robot travel by performing different laps of the environment (see Fig. 4, right), the states of the two internal neurons converge on a periodic limit cycle in which different states correspond to

different locations of the robot in the environment (Fig. 5).

As we mentioned above, the ability to generate this form of representation which allows the robot to solve its adaptive problem originates from the coupling between a simple robot's internal dynamics and a simple robot/body/environmental dynamics. The former dynamics is characterized by the fact that the state of the two internal neurons tends to move slowly toward different fixed point attractors, in the robot's internal dynamics, which correspond to different types of sensory states exemplified in Fig. 5. The latter dynamics originate from the fact that different types of sensory states last for different time durations and alternate with a given order while the robot moves in the environment. The interaction between these two dynamical processes leads to a transient dynamics of agents' internal state which moves slowly toward the current fixed point attractor without never fully reaching it (thus preserving information about previously experienced sensory states, the time duration of these states, and the order with which they have been experienced). The coupling

between the two dynamical processes originates from the fact that the free parameters which regulate the agent/environmental dynamics (e.g., the trajectory and the speed with which the robot moves in the environment) and the agent's internal dynamics (e.g., the direction and the speed with which the internal neurons change their state) have been coadapted and co-shaped during the adaptive process.

For related works which show how navigation and localization skills might emerge from the coupling between the agent's internal and external dynamics, see Tani and Fukumura (1997). For other works addressing other behavioral/cognitive capabilities, see Beer (2003) for what concerns categorization, Goldenberg et al. (2004) and Slocum et al. (2000) for what concerns selective attention, and Sugita and Tani (2005) for what concerns language and compositionality.

Behavior and Cognition as Phenomena with a Multilevel and Multi-scale Organization

Another fundamental feature that characterizes behavior is the fact that it is a multilayer system with different levels of organizations extending at different timescales (Baldassarre et al. 2006; Keijzer 2001). More precisely, as exemplified in Fig. 6, the behavior of an agent or of a group of agents involves both lower- and higher-level behaviors which extend for shorter or longer time spans, respectively. Lower-level behaviors arise from few agent/environmental interactions and short-term internal dynamical processes. Higher-level behaviors, instead, arise from the combination and interaction of lower-level behaviors and/or from long-term internal dynamical processes.


The multilevel and multi-scale organization of agents' behavior plays important roles: it is one of the factors which allow agents to produce functionally useful behavior without necessarily developing dedicated control mechanisms (Brooks 1991a, b; Nolfi 2005); it might favor the development of new behavioral and/or cognitive skills, thanks to the recombination and

reuse of preexisting capabilities (Marocco and Nolfi 2007); and it allows agents to generalize their skills in new task/environmental conditions (Nolfi 2005).

An exemplification of how the multilevel and multi-scale organization of behavior allows agents to generalize their skill in new environmental conditions is represented by the experiments carried out by Baldassarre et al. (2006) in which the authors evolved the control system of a group of robots assembled into a linear structure (Fig. 7) for the ability to move in a coordinated manner and for the ability to display a coordinated light-approaching behavior.


Each robot (Mondada et al. 2004) consists of a mobile base (chassis) and a main body (turret) that can rotate with respect to the chassis along the vertical axis. The chassis has two drive mechanisms that control the two corresponding tracks and toothed wheels. The turret has one gripper, which allows robots to assemble together and to grasp objects, and a motor controlling the rotation of the turret with respect to the chassis. Robots are provided with a traction sensor, placed at the turret-chassis junction, that detects the intensity and the direction of the force that the turret exerts on the chassis (along the plane orthogonal to the vertical axis) and light sensors. Given that the orientations of individual robots might vary and given that the target light might be out of sight, robots need to coordinate to choose a common direction of movement and to change their direction as soon as one or few robots start to detect a light gradient.

Evolved individuals show the ability to negotiate a common direction of movement and by approaching light targets as soon as a light gradient is detected. By testing evolved robots in different conditions, the authors observed that they are able to generalize their skills in new conditions and also to spontaneously produce new behaviors which have not been rewarded during the evolutionary process. More precisely, groups of assembled robots display a capacity to generalize their skills with respect to the number of robots which are assembled together and to the shape formed by the assembled robots. Moreover, when the evolved controllers are embodied



Embodied and Situated Agents, Adaptive Behavior in, Fig. 6 A schematic representation of multilevel and multi-scale organization of behavior. The behaviors represented in the *inner circles* represent elementary behaviors which arise from fine-grained interactions between the control system, the body, and the environment and which extend over limited time spans. The behaviors represented in the *external circles* represent


higher-level behaviors which arise from the combination and interaction between lower-level behaviors and which extend over longer time spans. The *arrows* which go from higher-level behavior toward lower levels indicate the fact that the behaviors currently exhibited by the agents later affect the lower-level behaviors and/or the fine-grained interaction between the constituting elements (agent's control system, agent's body, and the environment)



Embodied and Situated Agents, Adaptive Behavior in, Fig. 7 *Left*, four robots assembled into a linear structure. *Right*, a simulation of the robots shown in the *left part* of the figure

in eight robots assembled so to form a circular structure and situated in the maze environment shown in Fig. 8, the robots display an ability to collectively avoid obstacles, to rearrange their shape so to pass through narrow passages, and to explore the environment. The ability to display

all these behavioral skills allows the robots to reach the light target even in large maze environments, i.e., even in environmental conditions which are rather different from the conditions that they experienced during the training process (Fig. 8).



Embodied and Situated Agents, Adaptive Behavior in, Fig. 8 The behavior produced by eight robots assembled into a circular structure in a maze environment including walls and cylindrical objects (represented with gray lines and circles). The robots start in the central portion of the maze and reach the light target located in the bottom-left side of the environment (represented with an empty circle) by exhibiting a combination of coordinated-movement behaviors, collective-obstacle-avoidance, and collective-light-approaching behaviors. The irregular lines, that represent the trajectories of the individual robots, show how the shape of the assembled robots changes during motion by adapting to the local structure of the environment

By analyzing the behavior displayed by the evolved robots tested in the maze environment, a complex multilevel organization can be observed. The simpler behaviors that can be identified consist of low-level individual behaviors which extend over short time spans:

1. A *move-forward behavior*, which consists of the individuals' ability to move forward when the robot is coordinated with the rest of the team, is oriented toward the direction of the light gradient (if any) and does not collide with obstacles. This behavior results from the combination of (a) a control rule which produces a move-forward action when the perceived traction has a low intensity and when difference between the intensity of the light perceived on the left and the right side of the robot is low and (b) the sensory effects of the execution of the move-forward action selected mediated by the external environment which does not produce a variation of the state of the sensors until the conditions that should be satisfied to produce these behaviors hold.
2. A *conformistic behavior* which consists of the individuals' ability to conform its orientation with that of the rest of the team when the two orientations differ significantly. This behavior results from the combination of (a) a control rule that makes the robot turn toward the direction of the traction when its intensity is significant and (b) the sensory effects produced by the execution of this action mediated by the external environment that lead to a progressive reduction of the intensity of the traction until the orientation of the robot conforms with the orientation of the rest of the group.
3. A *phototaxis behavior* which consists of the individuals' ability to orient toward the direction of the light target. This behavior results from the combination of (a) a control rule that makes the robot turn toward the direction in which the intensity of the light gradient is higher and (b) the sensory effects produced by the execution of this action mediated by the external environment that lead to a progressive reduction of the difference in the light intensity detected on the two sides of the robot until the orientation of the robot conforms with the direction of the light gradient.
4. An *obstacle-avoidance behavior* which consists of the individuals' ability to change direction of motion when the execution of a motor action produced a collision with an obstacle. This behavior results from the combination of (a) the same control rule leading to behavior #2 which makes the robot turn toward the direction of the perceived traction (which in this case is caused by the collision with the obstacle, while in the case of behavior #2, it is caused by the forces exerted by the other assembled robots) and (b) the sensory effects produced by the execution of the turning action mediated by the external environment which makes the robot turn until collisions do not prevent anymore the execution of a move-forward behavior.

The combination and the interaction between these three behaviors produce the following higher-level collective behaviors that extend over a longer time span:

1. A *coordinated-motion behavior* which consists in the ability of the robots to negotiate a common direction of movement and to keep moving along such direction by compensating further misalignments originating during motion. This behavior emerges from the combination and the interaction of the conformistic behavior (which plays the main role when robots are misaligned) and the move-forward behavior (which plays the main role when robots are aligned).
2. A *coordinated-light-approaching behavior* which consists in the ability of the robots to coordinately move toward a light target. This behavior emerges from the combination of the conformistic and the move-forward and the phototaxis behaviors (which is triggered when the robots detect a light gradient). The relative importance of the three control rules which lead to the three corresponding behaviors depends both on the strength of the corresponding triggering condition (i.e., the extent of lack of traction forces, the intensity of traction forces, and the intensity of the light gradient, respectively) and on priority relations among behaviors (i.e., the fact that the conformistic behavior tends to play a stronger role than the phototaxis behavior).
3. A *coordinated-obstacle-avoidance behavior* which consists in the ability of the robots to coordinately turn to avoid nearby obstacles. This behavior arises as the result of the combination of the obstacle avoidance and the conformistic and the move-forward behaviors.

The combination and the interaction between these behaviors lead to the following higher-level collective behaviors that extend over still longer time spans:

1. A *collective-exploration behavior* which consists in the ability of the robots to visit different areas of the environment when the light

target cannot be detected. This behavior emerges from the combination of the coordinated-motion behavior and the coordinated-obstacle-avoidance behavior which ensures that the assembled robots can move in the environment without getting stuck and without entering into limit cycle trajectories.

2. A *shape-rearrangement behavior* which consists in the ability of the assembled robots to dynamically adapt their shape to the current structure of the environment so to pass through narrow passages especially when the passages to be negotiated are in the direction of the light gradient. This behavior emerges from the combination and the interaction between coordinated-motion and coordinated-light-approaching behaviors mediated by the effects produced by relative differences in motion between robots resulting from the execution of different motor actions and/or from differences in the collisions. The fact that the shape of the assembled robots adapt to the current environmental structure so to facilitate the overcoming of narrow passages can be explained by considering that collisions produce a modification of the shape which affects on particular the relative position of the colliding robots.

The combination and the interaction of all these behavior lead to a still higher-level behavior:

1. A *collective-navigation behavior* which consists in the ability of the assembled robots to navigate toward the light target by producing coordinated movements, exploring the environment, passing through narrow passages, and producing a coordinated-light-approaching behavior (Fig. 8).

This analysis illustrates two important mechanisms which explain the remarkable generalization abilities of these robots. The first mechanism consists in the fact that the control rules which regulate the interaction between the agents and the environment so to produce certain behavioral skills in certain environmental conditions will produce different but related behavioral skills in

other environmental conditions. In particular, the control rules which generate the behaviors #5 and #6 for which evolving robots have been evolved in an environment without obstacles also produce behavior #7 in an environment with obstacles. The second mechanism consists in the fact that the development of certain behaviors at a given level of organization which extend for a given time span will automatically lead to the exhibition of related higher-level behaviors extending at longer time spans which originate from the interactions from the former behaviors (even if these higher-level behaviors have not been rewarded during the adaptation process). In particular, the combination and the interaction of behaviors #5, #6, and #7 (which have been rewarded during the evolutionary process or which arise from the same control rules which lead to the generation of rewarded behaviors) automatically lead to the production of behaviors #8, #9, and #10 (which have not been rewarded). Obviously, there is no warranty that the new behaviors obtained as a result of these generalization processes will play useful functions. However, the fact that these behaviors are related to the other functional behavioral skills implies that the probabilities that these new behaviors will play useful functions are significant.

These generalization mechanisms can also be exploited by agents during their adaptive process to generate behavioral skills which play new functionalities and which emerge from the combination and the interaction between preexisting behavioral skills playing different functions. Indeed, by analyzing the evolutionary course of another series of experiments, we observed how the development of new behavioral capacities often creates the adaptive basis for the development of further skills that are produced by reusing and recombining previously developed skills (De Greef and Nolfi 2010).

On the Top-Down Effect from Higher to Lower Levels of Organization

In the previous sections, we have discussed how the interactions between the agents' body, the

agents' control system, and the environment lead to behavioral and cognitive skills and how such skills have a multilevel and multi-scale organization in which the interaction between lower-level skills leads to the emergence of higher-level skills. However, higher-level skills also affect lower-level skills up to the fine-grained interaction between the constituting elements (agents' body, agents' control system, and environment). More precisely, the behaviors, which originate from the interaction between the agent and the environment and from the interaction between lower-level behaviors, later affect the lower-level behaviors and the interaction from which they originate. These bidirectional influences between different levels of organization can lead to circular causality (Kelso 1995) where high-level processes act as independent entities which constraint the lower-level processes from which they originate.

One of the most important effects of these top-down influences is that the behavior exhibited by an agent constraint the type of sensory patterns that the agent will experience later on (i.e., constraint the fine-grained agent/environmental interactions which determine the behavior that will be later exhibited by the agent). Since the complexity of the problem faced by an agent depends on the sensory information experienced by the agent itself, these top-down influences can be exploited in order to turn hard problems into simple ones.

One neat demonstration of this type of phenomena is given by the experiments conducted by Nolfi and Marocco (2002) in which a simulated finger robot with six degree of freedom provided with sensors of its joint positions and with rough touch sensors is asked to discriminate between cubic and spherical objects varying in size. The problem is not trivial since, in general terms, the sensory patterns experienced by the robot do not provide clear regularities for discriminating between the two types of objects. However, the type of sensory states which are experienced by the agent also depends on the behavior previously exhibited by the agent itself – agents exhibiting different behaviors might face simpler or harder problems. By evolving the robots in simulation

for the ability to solve this problem and by analyzing the complexity of the problem faced by robots of successive generations, the authors observed that the evolved robot manages to solve their adaptive problem on the basis of simple control rules which allow the robot to approach the object and to move following the surface of the object from left to right, independently from the object shape. The exhibition of this behavior in interaction with objects characterized by a smooth or irregular surface (in the case of spherical or cubic objects, respectively) ensures that the same control rules lead to two types of behaviors depending on the type of the object. These behaviors consist in following the surface of the object and then moving away from the object in the case of spherical objects and in following the surface of the object by getting stuck in a corner in the case of cubic objects. The exhibition of these two behaviors allows the agent to experience rather different proprioceptor states as a consequence of having had interacted with spherical or cubic object which nicely encodes the regularities which are necessary to differentiate the two types of objects.

For other examples which show how adaptive agents can exploit the fact that behavioral and cognitive processes which arise from the interaction between lower-level behaviors or between the constituting elements later affect these lower-level processes, see Beer (2003), Nolfi (2002), and Scheier et al. (1998).

Adaptive Methods

In this section, we briefly review the methods through which artificial embodied and situated agents can develop their skill autonomously while they interact at different levels of organization with the environment and eventually with other agents. These methods are inspired by the adaptive process observed in nature: evolution, maturation, development, and learning.


We will focus in particular on self-organized adaptive methodologies in which the role of the experimenter/designer is reduced to the minimum and in which the agents are free to develop

their strategy to solve their adaptive problems within a large number of potentially alternative solutions. This choice is motivated by the following considerations:

1. These methods allow agents to identify the behavioral and cognitive skills which should be possessed, combined, and integrated so to solve the given problem. In other words, these methods can come up with effective ways of decomposing the overall required skill into a collection of simpler lower-level skills. Indeed, as we showed in the previous section, evolutionary adaptive techniques can discover ways of decomposing the high-level requested skill into lower-level behavioral and cognitive skills to find solutions which are effective and parsimonious, thanks to the exploitation of properties emerging from the interaction between lower-level processes and skills and thanks to the recruitment of previously developed skills for performing new functions. In other words, these methods release the designer from the burden of deciding how the overall skill should be divided into a set of simpler skills and how these skills should be integrated. More importantly, these methods can come up with solutions exploiting emergent properties which would be hard to design (Harvey 2000; Nolfi and Floreano 2000).
2. These methods allow agents to identify how a given behavioral and cognitive skill can be produced, i.e., the appropriate fine-grained characteristics of agents' body structure and control rules regulating the agent/environmental interaction. As for the previous aspect, the advantage of using adaptive techniques lies not only in the fact that the experimenter is released from the burden of designing the fine-grained characteristics of the agents but also in the fact that adaptation might prove more effective than human design due to the inability of an external observer to foresee the effects of a large number of nonlinear interactions occurring at different levels of organization.
3. These methods allow agents to adapt to variations of the task, of the environment, and of the social conditions.

Embodied and Situated Agents, Adaptive Behavior in, Fig. 9

A schematic representation of the evolutionary process. The stripes with black and white squares represent individual genotypes. The rectangular boxes indicate the genome of a population of a certain generation. The small robots placed inside the square on the right part of the figure represent a group of robots situated in an environment which interact with the environment and between themselves



Current approaches, in this respect, can be grouped into two families which will be illustrated in the following subsections and which include evolutionary robotics methods and developmental robotics methods.

Evolutionary Robotics Methods

Evolutionary robotics (Floreano et al. 2008; Nolfi and Floreano 2000) is a method which allows to create embodied and situated agents able to adapt to their task/environment autonomously through an adaptive process inspired by natural evolution (Holland 1975) and, eventually, through the combination of evolutionary, developmental, and learning processes.

The basic idea goes as follows (Fig. 9). An initial population of different artificial genotypes, each encoding the control system (and possibly the morphology) of an agent, is randomly created. Each genotype is translated into a corresponding phenotype (i.e., a corresponding agent) which is then left free to act (move, look around, manipulate the environment, etc.) while its performance (fitness) with respect to a given task is

automatically evaluated. In cases in which this methodology is applied to collective behaviors, agents are evaluated in groups which might be heterogeneous or homogeneous (i.e., might consist of agents who differ not with respect to their genetic and phenotypic characteristics). The fittest individuals (those having higher fitness) are allowed to reproduce by generating copies of their genotype with the addition of changes introduced by some genetic operators (e.g., mutations, exchange of genetic material). This process is repeated for a number of generations until an individual or a group of individuals is born which satisfies the performance level set by the user.

The process that determines how a genotype (i.e., typically a string of binary values) is turned into a corresponding phenotype (i.e., a robot with a given morphology and control system) might consist of a simple one-to-one mapping or of a complex developmental process. In the former case, many of the characteristics of the phenotypic individual (e.g., the shape of the body, the number and position of the sensors and of the actuators, and in some case the architecture of the neural controller) are predetermined and

fixed, and the genotype encodes a vector of free parameters (e.g., the connection weights of the neural controller Nolfi and Floreano 2000). In the latter case, the genotype encode a set of rules that determine how the body structure and the control system of the individual grow during an artificial developmental process. Through these types of indirect developmental mappings, most of the characteristics of the phenotypical robot can be encoded in the genotype and subjected to the evolutionary adaptive process (Nolfi and Floreano 2000; Pollack et al. 2001). Finally, in some cases, the adaptation process might involve both an evolutionary process that regulates how the characteristics of the robots vary phylogenetically (i.e., throughout generations) and a developmental/learning process which regulates how the characteristics of the robots vary ontogenetically (i.e., during the phase in which the robots act in the environment Nolfi and Floreano 1999).

Evolutionary methods can be used to allow agents to develop the requested behavioral and cognitive skills from scratch (i.e., starting from agents which do not have any behavioral or cognitive capability) or in an incremental manner (i.e., starting from pre-evolved robots which already have some behavioral capability which consists, for example, in the ability to solve a simplified version of the adaptive problem).

The fitness function which determines whether an individual will be reproduced or not might also include, in the addition to a component that scores the performance of the agent with respect to a given task, additional task-independent components. These additional components, in fact, can lead to the development of behavioral skills which are not necessarily functional but which can favor the development of functional skills later on (Prokopenko et al. 2006).

Evolutionary methods can allow agents to develop low-level behavioral and cognitive skills which have been previously identified by the designer/experimenter, which might later be combined and integrated in order to realize the high-level requested skill or directly to develop the high-level requested skill. In the former case, the adaptive process leads to the identification of the fine-grained features of the agent (e.g.,

number and type of sensors, body shape, architecture, and connection weights of the neural controller) which by interacting between themselves and with the environment will produce the required skill. In the latter case, the adaptive process leads to the identification of the lower-level skills (at different levels of organization) which are necessary to produce the required high-level skill, the identification of the way in which these lower-level skills should be combined and integrated and (as for the formed case) the identification of the fine-grained features of the agent which, in interaction with the physical and social environment, will produce the required behavioral or cognitive skills.

Developmental Robotics Methods

Developmental robotics (Asada et al. 2001; Brooks et al. 1998; Lungarella et al. 2003), also known as *epigenetic robotics*, is a method for developing embodied and situated agents that adapt to their task/environment autonomously through processes inspired by biological developmental and learning processes.

Evolutionary and developmental robotics methods share the same fundamental assumptions and also present differences for what concerns the way in which they are realized and the type of situations in which they are typically applied. For what concerns the former aspect, unlike evolutionary robotics methods which operate on long phylogenetic timescales, developmental methods typically operate on short ontogenetic timescales. For what concerns the latter aspects, unlike evolutionary methods which are usually used to develop behavioral and cognitive skills from scratch, developmental methods are typically adopted to model the development of complex developmental and cognitive skills from simpler preexisting skills which represent prerequisites for the development of the required skills.

At the present stage, developmental robotics does not consist of a well-defined methodology (Asada et al. 2001; Lungarella et al. 2003) but rather of a collection of approaches and methods

often addressing complementary aspects which hopefully would be integrated in a single methodology in the future. The sections below briefly summarize some of the most important methodological aspects of the developmental robotics approach.

The Incremental Nature of the Developmental Process

Development should be characterized as an incremental process in which preexisting structures and behavioral skills constitute important prerequisites and constraints for the development of more complex structures and behavioral skills and in which the complexity of the internal and external characteristics increases during development. One crucial aspect of developmental approach therefore consists in the identification of the initial characteristics and skills which should enable the bootstrapping of the developmental process: the layering of new skills on top of existing ones (Brooks et al. 1998; Metta et al. 2001; Scassellati 2001). Another important aspect consists in shaping the developmental process so to ensure that the progressive increase in the complexity of the task matches the current competency of the system and so to drive the developmental process toward the progressive acquisition of the skills which represent the prerequisites for further developments. The progressive increase in complexity might concern not only the complexity of the task or of the required skills but also the complexity of single component of the robot/environmental interaction, for example, the number of freeze/unfreeze degrees of freedom (Berthouze and Lungarella 2004).

The Social Nature of the Developmental Process

Development should involve social interaction with human subjects and with other developing robots. Social interactions (e.g., scaffolding, tutelage, mimicry, emulation, and imitation), in fact, play an important role not only for the

development of social skills (Breazeal 2003) but also as facilitators for the development of individual cognitive and behavioral skills (Tani et al. 2008). Moreover, other types of social interactions (i.e., alignment processes or social games) might lead to the development of cognitive and/or behavioral skills which are generated by a collection of individuals and which could not be developed by a single individual robot (Steels 2003).

Exploitation of the Interaction Between Concurrent Developmental Processes

Development should involve the exploitation of properties originating from the interaction and the integration of several co-occurring processes. Indeed, the codevelopment of different skills at the same time can favor the acquisition of the corresponding skills and of additional abilities arising from the combination and the integration of the developed skills. For example, the development of an ability to anticipate the sensory consequences of our own actions might facilitate the concurrent development of other skills such as categorical perception skills (Tani and Nolfi 1999). The development of an ability to pay attention to new situations (curiosity) and to look for new experiences after some time (boredom) might improve the learning of a given functional skill (Oudeyer et al. 2007; Schmidhuber 2006). The codevelopment of behavioral and linguistic skills might favor the acquisition of the corresponding skills and the development of semantic combinatoriality skills (Sugita and Tani 2005).

Discussion and Conclusion

In this entry, we described how artificial agents which are embodied and situated can develop behavioral and cognitive skills autonomously while interacting with their physical and social environment.

After having introduced the notion of embodiment and situatedness, we illustrated how the

behavioral and cognitive skills displayed by adaptive agents can be properly characterized as complex system with multilevel and multi-scale properties resulting from a large number of interaction at different levels of organization and involving both bottom-up processes (in which the interaction between elements at lower levels of organization leads to higher-level properties) and top-down processes (in which properties at a certain level of organization later affect lower-level properties or processes).

Finally, we briefly introduced the methods which can be used to synthesize adaptive embodied and situated agents.

The complex system nature of adaptive agents which are embodied and situated has important implications which constraint the organization of these systems and the dynamics of the adaptive process through which they develop their skills.

For what concerns the organization of these systems, it implies that agents' behavioral and/or cognitive skills (at any stage of the adaptive process) cannot be traced back to anyone of the three foundational elements (i.e., the body of the agents, the control system of the agents, and the environment) in isolation but should rather be characterized as properties which emerge from the interactions between these three elements and the interaction between behavioral and cognitive properties emerging from the former interactions at different levels of organizations. Moreover, it implies that complex behavioral or cognitive skills might emerge from the interaction between simple properties and processes.

For what concerns agents' adaptive process, it implies that the development of new complex skills does not necessarily require the development of new complex morphological features or new complex control mechanisms. Indeed, new complex skills might arise from the addition of new simple features or new simple control rules which, in interaction with the preexisting features and processes, might produce the required new behavioral or cognitive skills.

The study of adaptive behavior in artificial agents which has been reviewed in this entry has an important implication both from an engineering point of view (i.e., for progressing in our

ability to develop effective machines) and from a modeling point of view (i.e., for understanding the characteristics of biological organisms).

In particular, from an engineering point of view, progresses in our ability to develop adaptive embodied and situated agents can lead to development of machines playing useful functionalities.

From a modeling point of view, progresses in our ability to model and analyze artificial adaptive agents can improve our understanding of the general mechanisms behind animal and human intelligence. For example, the comprehension of the complex system nature of behavioral and cognitive skills illustrated in this entry can allow us to better define the notion of embodiment and situatedness which represent two foundational concepts in the study of natural and artificial intelligence. Indeed, although possessing a body and being in a physical environment certainly represent a prerequisite for considering an agent embodied and situated, a more useful definition of embodiment (or of truly embodiment) can be given in terms of the extent to which a given agent exploits its body characteristics to solve its adaptive problem (i.e., the extent to which its body structure is adapted to the problem to be solved, or, in other words, the extent to which its body performs morphological computation). Similarly, a more useful definition of situatedness (or truly situatedness) can be given in terms of the extent to which an agent exploits its interaction with the physical and social environment and the properties originating from this interaction to solve its adaptive problem. For the sake of clarity, we can refer to the former definition of the terms (i.e., possessing a physical body and being situated in a physical environment) as embodiment and situatedness in a weak sense and to the latter definition as embodiment and situatedness in a strong sense.

Bibliography

- Asada M, MacDorman K, Ishiguro H, Kuniyoshi Y (2001) Cognitive developmental robotics as a new paradigm for the design of humanoid robots. *Robot Auton Syst* 37:185–193
- Baldassarre G, Parisi D, Nolfi S (2006) Distributed coordination of simulated robots based on self-organisation. *Artif Life* 3(12):289–311

- Beer RD (1995) A dynamical systems perspective on agent-environment interaction. *Artif Intell* 72:173–215
- Beer RD (2003) The dynamics of active categorical perception in an evolved model agent. *Adapt Behav* 11:209–2
- Berthouze L, Lungarella M (2004) Motor skill acquisition under environmental perturbations: on the necessity of alternate freezing and freeing. *Adapt Behav* 1(1):47–63
- Bongard JC, Paul C (2001) Making evolution an offer it can't refuse: morphology and the extradimensional bypass. In: Keleman J, Sosik P (eds) *Proceedings of the sixth European conference on artificial life. Lecture notes in artificial intelligence*, vol 2159. Springer, Berlin
- Breazeal C (2003) Towards sociable robots. *Robot Auton Syst* 42(3–4):167–175
- Brooks RA (1991a) Intelligence without reason. In: Mylopoulos J, Reiter R (eds) *Proceedings of 12th international joint conference on artificial intelligence*. Morgan Kaufmann, San Mateo
- Brooks RA (1991b) Intelligence without reason. In: *Proceedings of 12th international joint conference on artificial intelligence*, Sydney, pp 569–595
- Brooks RA, Breazeal C, Irie R, Kemp C, Marjanovic M, Scassellati B, Williamson M (1998) Alternate essences of intelligence. In: *Proceedings of the fifteenth national conference on artificial intelligence (AAAI-98)*, Madison, pp 961–976
- Chiel HJ, Beer RD (1997) The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment. *Trends Neurosci* 20:553–557
- Clark A (1997) *Being there: putting brain, body and world together again*. MIT Press, Cambridge
- De Greef J, Nolfi S (2010) Evolution of implicit and explicit communication in a group of mobile robots. In: Nolfi S, Mirolli M (eds) *Evolution of communication and language in embodied agents*. Springer, Berlin
- Endo I, Yamasaki F, Maeno T, Kitano H (2002) A method for co-evolving morphology and walking patterns of biped humanoid robot. In: *Proceedings of the IEEE conference on robotics and automation*, Washington, DC
- Floreano D, Husband P, Nolfi S (2008) Evolutionary robotics. In: Siciliano B, Oussama K (eds) *Handbook of robotics*. Springer, Berlin
- Gigliotta O, Nolfi S (2008) On the coupling between agent internal and agent/environmental dynamics: development of spatial representations in evolving autonomous robots. *Adapt Behav* 16:148–165
- Goldenberg E, Garcowski J, Beer RD (2004) May we have your attention: analysis of a selective attention task. In: Schaal S, Ijspeert A, Billard A, Vijayakumar S, Hallam J, Meyer J-A (eds) *From animals to animats 8: proceedings of the eighth international conference on the simulation of adaptive behavior*. MIT Press, Cambridge
- Harvey I (2000) Robotics: philosophy of mind using a screwdriver. In: Gomi T (ed) *Evolutionary robotics: from intelligent robots to artificial life*, vol III. AAI Books, Ontario
- Holland J (1975) *Adaptation in natural and artificial systems*. University of Michigan Press, Ann Arbor
- Keijzer F (2001) *Representation and behavior*. MIT Press, London
- Kelso JAS (1995) *Dynamics patterns: the self-organization of brain and behaviour*. MIT Press, Cambridge
- Lungarella M, Metta G, Pfeifer R, Sandini G (2003) *Developmental robotics: a survey*. *Connect Sci* 15:151–190
- Marocco D, Nolfi S (2007) Emergence of communication in embodied agents evolved for the ability to solve a collective navigation problem. *Connect Sci* 19(1):53–74
- Massera G, Cangelosi A, Nolfi S (2007) Evolution of prehension ability in an anthropomorphic neurorobotic arm. *Front Neurorobot* 1(4):1–9
- McGeer T (1990) Passive walking with knees. In: *Proceedings of the IEEE conference on robotics and automation*, vol 2, pp 1640–1645
- Metta G, Sandini G, Natale L, Panerai F (2001) Development and Q30 robotics. In: *Proceedings of IEEE-RAS international conference on humanoid robots*, pp 33–42
- Mondada F, Franzi E, Ienne P (1993) Mobile robot miniaturisation: a tool for investigation in control algorithms. In: *Proceedings of the third international symposium on experimental robotics*. Kyoto
- Mondada F, Pettinaro G, Guigrard A, Kwee I, Floreano D, Denebourg J-L, Nolfi S, Gambardella LM, Dorigo M (2004) Swarm-bot: a new distributed robotic concept. *Auton Robot* 17(2–3):193–221
- Nolfi S (2002) Power and limits of reactive agents. *Neurol Comput* 49:119–145
- Nolfi S (2005) Behaviour as a complex adaptive system: on the role of self-organization in the development of individual and collective behaviour. *Complexus* 2(3–4):195–203
- Nolfi S, Floreano D (1999) Learning and evolution. *Auton Robot* 1:89–113
- Nolfi S, Floreano D (2000) *Evolutionary robotics: the biology, intelligence, and technology of self-organizing machines*. MIT Press/Bradford Books, Cambridge
- Nolfi S, Marocco D (2002) Active perception: a sensorimotor account of object categorization. In: Hallam B, Floreano D, Hallam J, Hayes G, Meyer J-A (eds) *From animals to animats 7, Proceedings of the VII international conference on simulation of adaptive behavior*. MIT Press, Cambridge, pp 266–271
- Oudeyer P-Y, Kaplan F, Hafner V (2007) Intrinsic motivation systems for autonomous mental development. *IEEE Trans Evol Comput* 11(2):265–286
- Pfeifer R, Bongard J (2007) *How the body shape the way we think*. MIT Press, Cambridge
- Pfeifer R, Iida F, Gómez G (2006) Morphological computation for adaptive behavior and cognition. *Int Congr Ser* 1291:22–29
- Pollack JB, Lipson H, Funes P, Hornby G (2001) Three generations of coevolutionary robotics. *Artif Life* 7:215–223

- Prokopenko M, Gerasimov V, Tanev I (2006) Evolving spatiotemporal coordination in a modular robotic system. In: Rocha LM, Yaeger LS, Bedau MA, Floreano D, Goldstone RL, Vespignani A (eds) *Artificial life X: proceedings of the tenth international conference on the simulation and synthesis of living systems*. MIT Press, Boston
- Scassellati B (2001) Foundations for a theory of mind for a humanoid robot. PhD thesis, Department of Electrical Engineering and Computer Science, MIT, Boston
- Scheier C, Pfeifer R, Kuniyoshi Y (1998) Embedded neural networks: exploiting constraints. *Neural Netw* 11:1551–1596
- Schmidhuber J (2006) Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. *Connect Sci* 18(2):173–187
- Schmitz A, Gómez G, Iida F, Pfeifer R (2007) On the robustness of simple speed control for a quadruped robot. In: *Proceeding of the international conference on morphological computation*, Venice
- Slocum AC, Downey DC, Beer RD (2000) Further experiments in the evolution of minimally cognitive behavior: from perceiving affordances to selective attention. In: Meyer J, Berthoz A, Floreano D, Roitblat H, Wilson S (eds) *From animals to animats 6. Proceedings of the sixth international conference on simulation of adaptive behavior*. MIT Press, Cambridge
- Steels L (2003) Evolving grounded communication for robots. *Trends Cogn Sci* 7(7):308–312
- Sugita Y, Tani J (2005) Learning semantic combinatoriality from the interaction between linguistic and behavioral processes. *Adapt Behav* 13(1):33–52
- Tani J, Fukumura N (1997) Self-organizing internal representation in learning of navigation: a physical experiment by the mobile robot Yamabico. *Neural Netw* 10(1):153–159
- Tani J, Nolfi S (1999) Learning to perceive the world as articulated: an approach for hierarchical learning in sensory-motor systems. *Neural Netw* 12:1131–1141
- Tani J, Nishimoto R, Namikawa J, Ito M (2008) Co-developmental learning between human and humanoid robot using a dynamic neural network model. *IEEE Trans Syst Man Cybern B Cybern* 38:1
- Varela FJ, Thompson E, Rosch E (1991) *The embodied mind: cognitive science and human experience*. MIT Press, Cambridge
- van Gelder TJ (1998) The dynamical hypothesis in cognitive science. *Behav Brain Sci* 21:615–628
- Vaughan E, Di Paolo EA, Harvey I (2004) The evolution of control and adaptation in a 3D powered passive dynamic walker. In: Pollack J, Bedau M, Husband P, Ikegami T, Watson R (eds) *Proceedings of the ninth international conference on the simulation and synthesis of living systems*. MIT Press, Cambridge

Interaction-Based Computing in Physics

Franco Bagnoli

Department Physics and Astronomy and CSDC,
University of Florence, Florence, Italy

Article Outline

Glossary

Definition

Introduction: Physics and Computers

From Trajectories to Statistics and Back

Artificial Worlds

Future Directions

Bibliography

Glossary

Correlation The correlation between two variables is the difference between the joint probability that the two variables take some values and the product of the two probabilities (which is the joint probability of two uncorrelated variables), summed over all possible values. In an extended system, it is expected that the correlation among parts diminishes with their distance, typically in an exponential manner.

Critical phenomenon A condition for which an extended system is correlated over extremely long distances.

Extended system A system composed by many parts connected by a network of interactions that may be regular (lattice) or irregular (graph).

Graph, lattice, tree A graph is set of nodes connected by links, oriented or not. If the graph is translationally invariant (it looks the same when changing nodes), it is called a (regular) lattice. A disordered lattice is a lattice with a fraction of removed links or nodes. An ordered set of nodes connected by links is called a path. A closed path not passing on

the same links is a loop. A cluster is a set of connected nodes. A graph can be composed by one cluster (a connected graph) or more than one (a disconnected graph). A tree is a connected graph without loops.

Mean field An approximate technique for computing the value of the observables of an extended system, neglecting correlations among parts. If necessary, the dynamics is first approximated by a stochastic process. In its simpler version, the probability of a state of the system is approximated by the product of the probability of each component, neglecting correlations. Since the state of two components that depend on a common “ancestor” (that interact with a common node) is in general not uncorrelated, and this situation corresponds to an interaction graph with loops, the simplest mean-field approximation consists in replacing the graph or the lattice of interactions with a tree.

Monte Carlo A method for producing stochastic trajectories in the state space designed in such a way that the time-averaged probability distribution is the desired one.

Nonlinear system A system composed by parts whose combined effects are different from the sum of the effects of each part.

Percolation The appearance of a “giant component” (a cluster that contains essentially all nodes or links) in a graph or a lattice, after adding or removing nodes or links. Below the percolation threshold, the graph is partitioned into disconnected clusters, none of which contains a substantial fraction of nodes or links, in the limit of an infinite number of nodes and links.

Probability distribution The probability of finding a system in a given state, for all the possible states.

State of a system A complete characterization of a system at a given time, assigning or measuring the positions, velocities, and other dynamical variables of all the elements (sometimes called a configuration). For completely discrete systems (cellular automata) of finite size, the state

of the system is just a set of integer numbers, and therefore the state space is numerable.

State space The set of all possible states of a system.

Trajectory A sequence of states of a system, labeled with the time, i.e., a path in the state space.

Definition

Theoretical physics investigation is based on building models of reality. Due to our limited cognitive capabilities, In order to *understand* a phenomenon, we need to represent it using a limited amount of symbols. Clearly, it is generally impossible to reproduce all aspects of a physical system with a simple model. However, even using simple *building blocks*, one can obtain systems whose behavior is quite complex. Therefore, most of the job of theoretical physics is that of identifying and investigating simple models that reproduce the main aspects of complex phenomena. Although some progress can be done analytically, most of models are also investigated by computer simulations.

Computers have changed the way a physical model is studied. Computers may be used to *calculate* the properties of a very complicated model representing a real system or to investigate *experimentally* what are the essential ingredients of a complex phenomenon. In order to carry out these explorations, several basic models have been developed, which are now used as building blocks for performing simulations and designing algorithms in many fields, from chemistry to engineering, from natural sciences to psychology. Rather than being derived from some fundamental law of physics, these blocks constitute *artificial worlds* still to be completely explored.

In this entry, we shall first present a pathway from Newton's laws to cellular automata and agent-based simulations, showing (some) computational approaches in classical physics. Then, we shall present some examples of *artificial worlds*.

Introduction: Physics and Computers

Some 60 years ago, shortly after the end of the Second World War, computers become available

to scientists. Actually, computers had already been used during the last years of the war for performing computations about a specific field of physics: the atomic bomb (Harlow and Metropolis 1983; Metropolis et al. 1980).

Up to then, the only available tool, except experiments, was paper and pencil. Starting with Newton and Leibnitz, humans discovered that continuous mathematics (i.e., differential and integral calculus) allowed to derive many consequences of a given hypothesis just by the manipulation of symbols. It seemed natural to express all quantities (e.g., time, space, mass) as continuous variables. Notice however that the idea of a continuous number is not at all "natural": one has to learn how to deal with it, while (small) integer numbers can be used and manipulated (added, subtracted) by illiterate humans and also by many animals. A point which is worth to be stressed is that any computation refers to a model of certain aspects of reality considered most important, while others are assumed to be not important.

Most of the investigations in physics only using paper and pencils are limited to almost-linear systems, or systems whose effective number of variables is quite small. On the other hand, most of the naturally occurring phenomena can be *successfully* modeled only using nonlinear elements. Therefore, a great deal of the precomputer physics is essentially linear physics, although astronomers (like other scientists) used to integrate numerically, by hand, the nonlinear equations of gravitation, in order to compute the trajectories of planets. This computation, however, was so cumbersome that no *playing* with trajectories was possible.

How can a computer solve physical problems? While analog computers have been used for integrating differential equations, the much more flexible digital computers are deterministic discrete systems. The way of working of a (serial) computer is that of a very fast automaton that manipulates data following a program.

In order to use computers as fast calculators, scientists ported and adapted existing numerical algorithms and developed new ones. This implied the use of techniques able to *approximate* the

computations of continuous mathematics using computer algebra. However, numbers in computers are not exactly the same as human numbers; in particular, they have finite (and varying) precision.

This intrinsic precision limit has deep consequences in the simulations of nonlinear systems, in particular of chaotic ones. Indeed, chaos was *numerically discovered* by Lorenz (Oestreicher 2007) after the observation that a simple approximation, a number that was retyped with fewer decimals, caused a macroscopic change in the trajectory under study.

With all their limits, computers can be fruitfully used *just* to speed up computations that *could* eventually be performed by humans. However, since the increase in velocity is of several order of magnitude, it becomes possible to include more and more details into the basic model of the phenomenon under investigation, well beyond what would be possible with an army of “human computers.” The idea of exploiting the *brute power* of fast computers has originated a fruitful line of investigation in *numerical physics* especially in the field of chemistry, biological molecules, and the structure of matter. The power of computers has allowed, for instance, to include quantum mechanical effects in the computation of the structure of biomolecules (Car and Parrinello 1985), and although these techniques may be targeted as “brute force,” the algorithms developed are actually quite sophisticated.

However, a completely different usage of computers is possible: instead of exploiting them for performing computations of models that already proved to approximate the reality, one can use computers as *experimental* apparatuses to investigate the patterns of *theoretical* models, which are generally nonlinear. This is what Lorenz did after having found the first example of chaos in computational physics. He started simplifying his equations in order to enucleate the minimal ingredients of what would be called the *butterfly effect*.

Much earlier than Lorenz, Fermi, Pasta, and Ulam (and the programmer Tsingou (Daxois et al. 2005)) used one on the very first available computers to investigate the basis of statistical mechanics: how energy distributes among the

oscillation modes of a chain of nonlinear oscillators (Fermi et al. 1955).

Also in this case the model is simplified at its maximum, in order to put into evidence what are the fundamental ingredients of the observed pattern and also to use all the available power of computers to increase the precision, the duration, and the size of the simulation.

This simplification is even more fruitful in the study of systems with many degrees of freedom that we may denote generically as *extended systems*. We humans are not prepared to manipulate more than a few symbols at once. So, unless there is a way of grouping together many parts (using averages, like, for instance, when considering the pressure of a gas as an average over extremely many particle collisions), we are in difficulties in *understanding* such systems. They may nevertheless be studied performing *experiments* on computers. Again, the idea is that of simplifying at most the original model, in order to isolate the fundamental ingredients of the observed behavior. It is therefore natural to explore systems whose *physics* is different from the usual one. These *artificial worlds* are preferably formulated in discrete terms, more suitable to be implemented in computers (see section “Artificial Worlds”).

This line of investigation is of growing interest today: since modern computers can easily simulate thousands or millions of *elementary agents*, it may be possible to design *artificial worlds* in which *artificial people* behave similarly to real humans. The rules of these worlds are not obtained from the *basic laws* of the real one, since no computer can at present simulate the behavior of all the elements of even a small portion of matter. These rules are designed so to behave similarly to the system under investigation and to be easily implemented in digital terms. There are two main motivations (or hopes): to be able to understand real complex dynamics by studying simplified models and to be so lucky to discover that a finely tuned model is able to reproduce (or forecast) the behavior of its real counterpart.

This is so promising that many scientists are performing experiments on these artificial worlds, in order to extract their principal characteristics, to be subsequently analyzed using paper and pencil!

In the following, we shall try to elucidate some aspects of the interplay between computer and physics. In section “[From Trajectories to Statistics and Back](#),” we shall illustrate possible logic pathways (in classical mechanics) leading from Newton’s equations to research fields that use computers as investigative tool, like agent-based investigations of human societies. In section “[Artificial Worlds](#),” we shall try to present succinctly some example of artificial systems that are still active research topics in theoretical physics.

From Trajectories to Statistics and Back

The outline of this section is the following. Physics is often denoted the most *fundamental* science, and one may think that given powerful enough computers, one should be able to reconstruct any experimental situation simply implementing the fundamental laws of physics. I would like to show that any investigation is based on models, requiring approximations and additional assumptions, and that any change of scale implies a change of model. However, an important suggestion from physics is that similar models can be used to interpret different situations, and therefore, the associated computational techniques can be *reused* in different contexts. We shall follow this line from Newton’s equations to agent-based modeling.

Let us assume that a working model of the reality can be built using a set of dynamical equations, for instance, those of classical mechanics. We shall consider the model of a system formed by many particles, like a solid or a fluid. The state of the resulting system can be represented as a point in a high-dimensional space, since it is given by all coordinates and velocities of all particles. The evolution of the system is a trajectory in such space. Clearly, the visualization and the investigation of such a problem is challenging, even using powerful computers.

Moreover, even if we were able to compute many trajectories (in order to have an idea of fluctuations), this would not imply that we have *understood* the problem. Let us consider, for instance, the meteorology: one is interested in

the probability of rain or in the expected wind velocity, not in forecasting the trajectories of all molecules of air. Similarly, in psychology, one is interested in the expected behavior of an individual, not in computing the activity of all neurons in his brain.

Since physics is the oldest discipline that has been *quantified* into equations, it may be illuminating to follow some of the paths followed by researchers to *reduce* the complexity of a high-dimensional problem to something more manageable or at least simpler to be simulated on a computer.

In particular, we shall see that many approaches consist in *projecting* the original space onto a limited number of dimensions, corresponding to the observables that vary in a slow and smooth way and assuming that the rest of the dynamics is eventually approximated by “noise.” The noise can be so small, compared to the macroscopic observables that it can be neglected. In such cases, one has a deterministic, low-dimensional dynamical system, for instance, the usual models for rigid bodies, planets, etc. In the opposite case, the resulting system is stochastic, and one is interested in computing the average values of observables, over the probability distribution of the projected system.

However, the computation of the probability distribution may be hard, and so one seeks to find a way of producing *artificial* trajectories, in the projected space, designed in such a way that their probability distribution is the desired one. So doing, the problem reduces to the computation of the time-averaged values of *slow* observables. For the rest of this section, please make reference to Fig. 1.

Newton’s Laws

The success of Newton in describing the motion of one body, subjected to a static field force (say, gravitational motion of one planet, oscillation of a body attached to a spring, motion of the pendulum, etc.), clearly proved the validity of his approach and also the validity of using simple models for dealing with natural phenomena. Indeed, the representation of a body as a point mass, the idea of massless springs and strings,

perturbation of the parameters or of the initial conditions with large variations of its trajectory. This *sensibility* to variation implies the impossibility of predicting its behavior for long times, unless one is content with a probabilistic description.

The reduction from a high-dimensional system to a small number of equation can in general be considered the result of a projection operation, similar to studying the dynamics of couple of dancers only looking at his shadow. The problem is of course that of finding the right projection, for instance, that able to show that the couple is formed by two distinct bodies that only occasionally separate.

High Dimensionality

In many cases, the *projection* operation results in a system still composed by many parts. For instance, models of nonequilibrium fluids neglect to consider the movement of the individual molecules, but still one has to deal with the values of the pressure, density, and velocity in all points. In these cases, one is left with a high-dimensional problem. Assuming that the *noise* of the projected dimensions can be neglected, one can either write down a large number of coupled equation (e.g., in modeling the vibration of a crystal) or use a continuous approach and describe the system by means of partial differential equations (e.g., the model of a fluid).

Linear Systems

In general, high- and low-dimensional approaches can be systematically developed (with paper and pencil) only in the linear approximation. Let us illustrate this point for the case of coupled differential equation: if the system is linear, one can write the equations using matrices and vectors. One can in principle find a (linear) transformation of variables that make the system diagonal, i.e., that reduces the problem to a set of *uncoupled* equations. At this point, one is left with (many) one-dimensional independent problems. Clearly, there may be mathematical difficulties, but the path is clear. A similar approach (for instance, using Fourier transforms) can be used also for dealing with partial differential equations.

The variables that result from such operations are called normal modes, because they behave independently one from the other (i.e., they correspond to orthogonal or normal directions in the state space). For instance, the linear model of a vibrating string (with fixed ends) predicts that any pattern can be described as a superposition of “modes,” which are the standing oscillations with zero, one, two, etc. nodes (the harmonics).

However, linear systems behave in a somewhat strange way, from the point of view of thermal physics. Let us consider, for instance, the system composed by two *uncoupled* oscillators. It is clear that if we excite one oscillator with any amount of energy, it will remain confined to that subsystem. With normal modes, the effect is the same: any amount of energy communicated to a normal mode remains confined to that mode, if the system is completely linear. In other words, the system never forgets its initial conditions.

On the contrary, the long-time behavior of *normal* systems does not depend strongly on the initial conditions. One example is the *timbre* or “sound color” of an object. It is given by the simultaneous oscillations on many frequencies, but in general an object emits its characteristic sound regardless of how exactly is excited. This would not be true for linear systems.

Since the *distribution* of energy to all available *modes* is one of the assumptions of equilibrium statistical mechanics, which allows us to *understand* the usual behavior of matter, we arrived at an unpleasant situation: linear systems, which are so “easy” to be studied, cannot be used to ground statistical physics on mechanics.

Molecular Dynamics

Nowadays, we have computers at our disposal, and therefore we can simulate systems composed by many parts with complex interactions. One can do this simply discretizing Newton’s equations of motion so that a digital computer can approximately integrate the set of coupled differential equations. This technique is sometimes called *molecular dynamics*.

One is generally interested in computing *macroscopic* quantities. These are defined as *averages* of some function of the microscopic variables (positions, velocities, accelerations, etc.) of the

system. A measurement on a system implies an average, over a finite interval of time and over a large number of elementary components (say, atoms, molecules, etc.) of some quantity that depends on the microscopic state of that portion of the body.

Chaos and Probability

It was realized by Poincaré (with paper and pencil) and Lorenz (with one of the very first computers) that also very few (three) coupled differential equations with nonlinear interactions may give origin to complex (chaotic) behavior. In a chaotic system, a small uncertainty amplifies exponentially in time, making forecasting difficult. However, chaos may also be simple: the equations describing the trajectory of dice are almost surely chaotic, but in this case the chaos is so strong that the tiniest perturbation or uncertainty in the initial conditions will cause in a very small amount of time a complete variation of the trajectory. Our experience says that the process is well approximated by a probabilistic description. Therefore, chaos is one possible way of introducing probability in dynamics.

Chaotic behavior may be obtained in simpler models, called maps, that evolve in discrete time steps. As May (1976) showed, a simple map with a quadratic nonlinearity (logistic map) may be chaotic. One can also model a system using coupled maps instead of a system of coupled differential equations (Kaneko 1985). And indeed, when a continuous system is simulated on a computer, it is always represented as an array of coupled maps.

Computers are also used to simulate stochastic systems, which at first may seem quite strange: how can a deterministic system, like a computer, generate random data? Although there is the possibility of using unpredictable events (like network data), amplifying thermal noise or even use quantum-mechanical events for generating true random numbers, in most of cases the pseudo-random number generators just use a chaotic map, properly initialized (for instance, using a true random number).

Discretization

Many extended systems (for instance, fluids) are described by means of partial differential equations, which are however quite difficult to be

handled by computers. When possible, it is more efficient to reduce the original system to a (small) number of ordinary differential equations that are then converted to discrete-time maps.

There is a therefore a progression of discretization from partial differential equations, coupled differential equations, coupled map lattices: from systems that are continuous in space, in time, and in the dynamical variables to systems that are discrete in time and space and continuous only in the dynamical variables. The further logic step is that of studying completely discrete systems, called *cellular automata*.

Cellular automata show a wide variety of different phenomenology. They can be considered mathematical tools or used to model reality. In many cases, the resulting phenomenological models follow probabilistic rules, but it is also possible to use cellular automata as *building blocks*. For instance, it is possible to simulate the behavior of a hydrodynamic system by means of a completely discrete model, called cellular automata lattice gas (Frisch et al. 1986; Hardy et al. 1973).

Statistics

The investigation of chaotic extended systems proceeds generally using a statistical approach. The idea is the following: any system contains a certain degree of nonlinearity that couples otherwise independent normal modes. Therefore, (one hopes that) the initial condition is not too important for the asymptotic regime. If moreover one assumes that the motion is so chaotic that any trajectory spans the available space in a *characteristic* way (again, not depending on the initial conditions), we can use statistics to derive the characteristic probability distribution: the probability of finding the system in a given portion of the available space is proportional to the time that the system spends in that region. See also the paragraph on equilibrium.

Random Walks

Another approach is that of focusing on a small part of a system, for instance, a single particle. The rest of the system is approximated by *noise*. This method was applied, for instance, by Einstein in the development of the simplest theory of Brownian motion, the random walk (Haw 2005).

In random walks, each step of the target particle is independent on previous steps, due to collisions with the rest of particles. Collisions, moreover, are supposed to be uncorrelated. A more sophisticated approximation consists in keeping some aspects of motion, for instance, the influence of inertia or of external forces, still approximating the rest of the world by noise (which may contain a certain degree of correlation). This is known as the Langevin approach, which includes the random walk as the simplest case. Langevin equations are stochastic differential equations.

The essence of this method relies in the assumption that the behavior of the various parts of the systems is uncorrelated. This assumption is vital also for other types of approximations that will be illustrated in the following. Notice that in the statistical mechanics approach, this assumption is not required.

In the Langevin formulation, by averaging over many *independent* realizations of the process (which in general is not the same of averaging over many particles that *simultaneously* move, due, for instance, to excluded volumes), one obtains the evolution equation of the probability of finding a particle in a given portion of space. This is the Kolmogorov integral-differential equation that in many case can be simplified, giving a differential (Fokker-Plank) equation. The diffusion equation is just the simplest case (Gardiner 1994; van Kampen 1992).

It is worth noticing that a similar formulation may be developed for quantum systems: the Feynman path-integral approach is essentially a Langevin formulation, and the Schroedinger equation is the corresponding Fokker-Plank equation.

Random walks and stochastic differential equations find many applications in economics, mainly in stock market simulations. In these cases, one is not interested in the average behavior of the market, but rather in computing nonlinear quantities over trajectories (time series of good values, fluctuations, etc.).

Time-Series Data

In practice, a model is never derived *ab initio*, by projecting the dynamics of all the microscopic components onto a limited number of dimensions,

but is constructed heuristically from observations of the behavior of a real system.

It is therefore crucial to investigate how observations are made, i.e., the analysis of a series of time measurements. In particular, a good exercise is that of simulating a dynamical or stochastic system, analyze the resulting time-series data of a given observable, and see if one is able to reconstruct from it the relationships or the equations ruling the time evolution.

Let us consider the experimental study of a chaotic, low-dimensional system. The measurements on this system give a time series of values that we assume discrete (which is actually the case considering experimental errors). Therefore, the output of our experiment is a series of symbols or numbers, a time series. Let us assume that the system is stationary, i.e., that the sequence is statistically homogeneous in time. If the system is not extremely chaotic, symbols in the sequence are correlated, and one can derive the probability of observing single symbols, couples of symbols, triples of symbols, etc. There is a hierarchy in these probabilities, since the knowledge of the distribution of triples allows the computation of the distribution of couples, and so on.

It can be shown that the knowledge of the probability distribution of the infinite sequence is equivalent to the complete knowledge of the dynamics. However, this would correspond to performing an infinite number of experiments, for all possible initial conditions.

The usual investigation scheme assumes that correlations vanish beyond a certain distance, which is equivalent to assume that the probability of observing sequences longer than that distance factorize. Therefore, one tries to model the evolution of the system by a probabilistic dynamics of symbols as shown in section “[Probabilistic Cellular Automata](#).” Time-series data analysis can therefore be considered as the main experimental motivation in developing probabilistic discrete models. This can be done heuristically comparing results with observations a posteriori or trying to extract the rules directly from data, like in the Markov approach.

Markov Approximation

The Markov approach, either continuous or discrete, also assumes that the memory of the system

vanishes after a certain time interval, i.e., that the correlations in time series decay exponentially. In discrete terms, one tries to describe the process under study as an automaton, with given transition probabilities. The main problem is, given a sequence of symbols, what is the simplest automaton (hidden Markov chains (Rabiner 1989)) that can generate that sequence with maximum “predictability,” i.e., with transition probabilities that are nearest to zero or one? Again, it is possible to derive a completely deterministic automaton, but in general it has a number of states equivalent to the length of the time series, so it is not generalizable and has no predictability (see also section “Probabilistic Cellular Automata”). On the contrary, an automaton with a very small number of nodes will have typically intermediate transition probabilities, so predictability is again low (essentially equivalent to random extraction). Therefore, the good model is the result of an optimization problem that can be studied using, for instance, Monte Carlo techniques.

Mean Field

Finally, from the probabilities one can compute averages of observables, fluctuations, and other quantities called *moments* of the distribution. Actually, the knowledge of all moments is equivalent to the knowledge of the whole distribution. Therefore, another approach is that of relating moments at different times or different locations, truncating the recurrences at a certain level. The roughest approximation is that of truncating the relations at the level of averages, i.e., the *mean-field approach*. It appears so natural that it is often used without realizing the implications of the approximations. For instances, chemical equations are essentially mean-field approximations of a complex phenomenon.

Boltzmann Equation

Another similar approach is that of dividing an extended system into zones and assume that the behavior of the system in each zone is well described by a probability distribution. By disregarding correlations with other zones, one obtains the Boltzmann equation, with which many transport phenomena may be studied well beyond

the elementary kinetic theory. The Boltzmann equation can also be obtained from the truncation of a hierarchy of equations (BBGKY hierarchy) relating multiparticle probability distributions. Therefore, the Boltzmann equation is similar in spirit to a mean-field analysis.

Equilibrium

One of the biggest successes of the stochastic approach is *equilibrium statistical mechanics*. The main ingredient of this approach is that of minimum information, which, in other words, corresponds to the assumption: *what is not known is not harmful*. By supposing that at equilibrium the probability distribution of the systems maximizes the information entropy (corresponding to a minimum of information on the system), one is capable of deriving the probability distribution itself and therefore the expected values of observables (ensemble averages, see section “Ising Model”). In this way, using an explicit model, one is capable to compute the value of the parameters that appear in thermodynamics. If it were possible to show that the maximum entropy state is actually the state originated by the dynamics of a mechanical (or quantum) system, one could ground thermodynamics on mechanics. This is a long-investigated subject, dating back to Boltzmann, which is however not yet clarified. The main drawback in the derivations is about ergodicity. Roughly speaking, a system is called *ergodic* if the infinite-time average of an observable over a trajectory coincides with its average over a snapshot of infinitely many replicas. For a system with fixed energy and no other conserved quantities, a sufficient condition is that a generic trajectory passes *near* all points of the accessible phase space. However, most systems whose behavior is “experimentally” well approximated by statistical mechanics are not ergodic. Moreover, another ingredient, the capability of quickly *forgetting* the information about initial conditions appears to be required; otherwise trajectories are strongly correlated, and averages over different trajectories cannot be “mixed” together. This capability is strongly connected to the chaoticity or *unpredictability* of extended systems, but unfortunately these ingredients make analytic approaches quite hard.

An alternative approach, due to Jaynes (1957), is much more pragmatic. In essence, it says: let design a model with the ingredients that one thinks are important, and assume that what is not in the model does not affect its statistical properties. Compute the distribution that maximizes the entropy with the given constraints. Then, compare the results (averages of observables) with experiments (possibly, numerical ones). If they agree, one has captured the essence of the problem; otherwise one has to include some other ingredients and repeat the procedure. Clearly, this approach is much more general than the “dynamical” one, not considering trajectory or making assumptions about the energy, which is simply seen as a constraint. But physicists would be much more satisfied by a *microscopic* derivation of statistical physics.

In spite of this lack of strong basis, the statistical mechanics approach is quite powerful, especially for systems that can be reduced to the case of *almost* independent elements. In this situation, the probability distribution of the system (the *partition function*) factorizes, and many computations may be performed by hand. Notice however that this behavior is in strong contrast to that of truly linear systems: the “almost” attribute indicates that actually the elements interact and therefore share the same “temperature.”

Monte Carlo

The Monte Carlo technique was invented for computing, with the aid of a computer, thermal averages of observables of physical systems at equilibrium. Since then, this term is often used to denote the technique of computing the average values of observables of a stochastic system by computing the time-average values over *artificial* trajectories.

In equilibrium statistical physics, one is faced by the problem of computing averages of observables over the probability distribution of the system, and since the phase space is very high dimensional, this is in general not an easy task: one cannot simply draw *random* configurations, because in general they are so different from those *typical* of the given value of the temperature, that their statistical weight is marginal. And one does

not want to revert to the original, still-more-highly dimensional dynamical system, which typically requires powerful computers just to be followed for tiny time intervals.

First of all, one can divide (*separate*) the model into almost independent subsystems that, due to the small residual interactions (the “almost” independency), are at the same temperature. In the typical example of a gas, the velocity components appear into the formula of energy as additive terms, i.e., they do not interact with themselves or with other variables. Therefore, they can be studied separately giving the *Maxwell distribution* of velocities. The positions of molecules, however, are linked by the potential energy (except in the case of an ideal gas), and so the hard part of the computation is that of generating configurations. Secondly, statistical mechanics guarantees that the asymptotic probability distribution does not depend on the details of dynamics. Therefore, one is free to look for the simplest dynamics still compatible with constraints. The Monte Carlo computation is just a set of recipes for generating such trajectories. In many problems, this approach allows to reduce the (computational) complexity of the problem of several orders of magnitude, allowing to generate *artificial* trajectories that span statistically significant configuration with small computational effort. In parallel with the generation of the trajectory, one can compute the value of several observables and perform statistical analysis on them, in particular the computation of time averages and fluctuations.

By extension, the same terms *Monte Carlo* is used for the technique of generating sequences of states (trajectories) given the transition probabilities, and computing averages of observables on trajectories, instead of over the probability distribution.

Stochastic Optimization

One of the most interesting applications of Monte Carlo simulations concerns stochastic optimization via *simulated annealing*. The idea is that of exploiting an analogy between the status of a system (and its energy) and the coding of a particular procedure with corresponding cost function. The goal is that of finding the *best* solution,

i.e., the global minimum of the energy given the constraints. *Easy* systems have a smooth energy landscape, shaped like a funnel, so that usual techniques like that of always choosing the displacements that locally lowers the energy (gradient descent) are successful. However, when the energy landscape is corrugated, there are many local minima where algorithms like gradient descent tend to get trapped. Methods from statistical mechanics (Monte Carlo), on the contrary, are targeted to generating trajectories that quickly explore the *relevant* parts of the state space, i.e., those that correspond to the largest contributions to the probability distribution that depends on the temperature, an *external* or *control* parameter. If the temperature is high, the probability distribution is broad and the generated trajectory does not “see” the minima of energy that are below the temperature, i.e., it can jump over and off the local minima.

By lowering the temperature, the probability distribution of system obeying statistical mechanics concentrates around minima of energy, and the Monte Carlo trajectory does the same. The energy (or cost) function of not extremely complex problems is shaped in such a way that the global optimum belongs to a broad valley, so that this lowering of the temperature increases the probability of finding it.

Therefore, a sufficiently slow annealing should furnish the desired global minimum. Moreover, it is possible to *convert* constraints into energy terms, which is quite convenient since for many problems, it is difficult to generate configurations that satisfy the constraints. Let us think, for instance, to the problem of generating a school timetable, keeping into consideration that lessons should not last more than three consecutive hours, that a teacher or students cannot stay in two classes at the same time, and that a teacher is not available on Monday, another prefers the first hours, etc. It is difficult to generate generic configurations that obey all constraints, while it is easy to formulate a Monte Carlo algorithm that generates arbitrary configurations, weighting them with a factor that depends on how many constraints are violated. At high temperature, constraints do not forbid the exploration of the state

space and therefore to try *creative* solutions. At low temperature, constraints become important. At zero temperature, the configurations with lower energy are those that satisfy all constraints, if possible, or at least the largest part of them.

In recent years, physics have dealt with extremely complex problems (e.g., spin glasses (Dotsenko 1994; Mezard et al. 1987)), in which the energy landscape is extremely rough. Special techniques based on nonmonotonous “walks” on temperature have been developed (simulated tempering (Marinari and Parisi 1992)).

Critical Phenomena

One of the most investigated topics of statistical mechanics concerns phase transitions. This is a fascinating subject: in the vicinity of a continuous phase transitions, correlation lengths diverge, and the system behave collectively, in a way which is largely independent of the details of the model. This universal behavior allows the use of extremely simplified models that therefore can be massively simulated.

The *philosophy* of statistical mechanics may be *exported* to nonequilibrium systems: systems with absorbing states (that correspond to infinitely negative energy), driven systems (live living ones), strongly frustrated systems (that actually never reach equilibrium), etc. In general, one defines these systems in terms of transition probabilities, not in terms of energy. Therefore, one cannot invoke a maximum entropy principle, and the results are less general.

However, many systems exhibit behavior reminiscent of equilibrium systems, and the same language can be used: phase transitions, correlations, susceptibilities, etc. These characteristics, common to many different models, are sometimes referred as *emergent features*.

One of the most famous problems in this field is percolation: the formations of giant clusters in systems described by a local aggregation dynamics, for instance, adding links to a set of nodes. This *basic* model has been used to describe an incredibly large range of phenomena (Stauffer and Aharony 1994).

Equilibrium and nonequilibrium phase transitions occur for a well-defined value of a control

parameter. However, in nature one often observes phenomena whose characteristics resemble that of a system near a phase transition, critical dynamics, without any *fine-tuned* parameter. For such system, the term self-organized criticality has been coined (Bak et al. 1987), and they are the subject of active researches.

Networks

A recent *extension* of statistical physics is the theory of networks. Networks in physics are often regular, like the lattice of a crystal, or only slightly randomized. Characteristics of these networks are the fixed (or slightly dispersed around the mean) number of connections per node, the high probability of having connected neighbors (number of “triangles”), and the large time needed to cross the network. The opposite of a regular network is a random graph, which, for the same number of connections, exhibits low number of triangles and short crossing time. The statistical and dynamical properties of systems whose connection are regular or random are generally quite different.

Watts and Strogatz (1998) argued that *social networks* are never completely regular. They showed that the simple random rewiring of a small number of links in a regular network may induce the *small world* effect: local properties, like the number of triangles, are not affected, but large-distance ones, like the crossing time, quickly became similar to that of random graphs. Also the statistical and dynamical properties of models defined over a rewired network are generally similar to those correlated to random graphs.

After this finding, many social networks were studied, and they revealed a yet different structure: instead of having a well-defined connectivity, many of them present a few highly connected “hubs” and a lot of poorly connected “leafs.” The distribution of connectivity is often shaped as a power law (or similar (Newman 2005)), without a well-defined mean (scale-free networks (Albert and Barabasi 2002)). Many of phenomenological models are presently reexamined in order to investigate their behavior over such networks. Moreover, scale-free networks cannot be *laid down*; they need to be *grown* following a procedure, similar in this to fractals. It is natural

therefore to include such a procedure in the model, so that they not only evolve *over* the networks but also evolve *the* network (Boccaletti et al. 2006).

Agents

Many of the described tools are used in the so-called agent-based modeling. The idea is that of exploiting the powerful capabilities of present computers to simulate directly a large number of *agents* that interact among them. Traditional investigations of *complex systems*, like crowds, flocks, traffic, and urban models, have been performed using *homogeneous* representation: partial differential equations (i.e., mean field), Markov equations, cellular automata, etc. In such an approach, it is supposed that each agent type is present in many identical copies, and therefore they are simulated as *macrovariables* (cellular automata) or aggregated like independent random walkers in the diffusion equation. But live elements (cells, organisms) do not behave in such a way: they are often individually unique, carry information about their own past history, and so on. With computers, we are now in the position of simulating large assemblies of individuals, possibly geographically located, like, for instance, humans in an urban simulation.

One of the advantages of such approach is that of offering the possibility of measuring quantities that are inaccessible to field researchers and also playing with different scenarios. The disadvantages are the proliferation of parameters that are often beyond experimental confirmation.

Artificial Worlds

A somewhat alternative approach to that of *traditional* computational physics is that of studying an artificial model, built with little or no direct connection with reality, trying to include only those aspects that are considered relevant. The goal is to be able to find the *simplest* system still able to exhibit the relevant features of the phenomena under investigation. The resulting models, though not directly applicable to the interpretation of experiments, may serve as interpretative tools in many different situations. For instance, the Ising

model was developed in the context of the modeling of magnetic systems but has been applied to opinion formation, social simulations, etc.

Ising Model

The Ising (better: Lenz-Ising) model is probably one of the most known models in statistical physics. Its history (Niss 2005) is particularly illuminating in this context, even if it took place well before the advent of computers in physics. It is also a model for which the Monte Carlo and simulated annealing techniques are readily applied.

Let us first illustrate schematically the model. I shall present the traditional version, with the terminology that arises from the physics of magnetic systems. However, it is an interesting exercise to reformulate it in the context, for instance, of opinion formation. Let us simply replace “spin up/down” with “opinion A/B,” “magnetization” with “average opinion,” “coupling” with “exchange of ideas,” “external magnetic field” with “propaganda,” and so on.

The Ising model is defined on a lattice that can be in one, two, or more dimensions or even on a disordered graph. We shall locate a cell with an index i , corresponding to the set of spatial coordinates for a regular lattice or a label for a graph. The dynamical variable x_i for each cell is just a binary digit, traditionally named “spin” that takes the values ± 1 . We shall indicate the whole configuration as $\mathbf{x}$. Therefore, a lattice with N cells has 2^N distinct configurations. Each configuration $\mathbf{x}$ has an associated energy

$$E(\mathbf{x}) = -\sum_i (H + h_i) x_i,$$

where H represents the external magnetic field and h_i is a local magnetic field, generated by neighboring spins, $h_i = \sum_j J_{ij} x_j$. The coupling J_{ij} for the original Lenz-Ising model is one if i and j are nearest neighbors, and zero otherwise.

The maximum-entropy principle (Jaynes 1957) gives the probability distribution

$$P(\mathbf{x}) = (1/Z) \exp(-E(\mathbf{x})/T)$$

from which averages can be computed. The parameter T is the temperature, and Z , the “partition function,” is the normalization factor of the distribution.

The quantity $E(\mathbf{x})$ can be thought as a *landscape*, with low-energy configurations corresponding to valley and high-energy ones to peaks. The distribution $P(\mathbf{x})$ can be interpreted as the density of a gas, each “particle” corresponding to a possible realization (a replica) of the system. This *gas* concentrates in the valleys for low temperatures and diffuses if the temperature is increased. The temperature is related to the average level of the gas.

In the absence of the local field ($J = 0$), the energy is minimized if each x_i is *aligned* (same sign) with H . This ordering is counteracted by thermal noise. In this case, it is quite easy to obtain the average magnetization per spin (order parameter)

$$\langle x \rangle = \tanh(H/T),$$

which is a plausible behavior for a paramagnet. A ferromagnet however presents hysteresis; i.e., it may maintain for long times (metastability) a pre-existing magnetization opposed to the external magnetic field.

With coupling turned on ($J > 0$), it may happen that the local field is strong enough to “resist” H , i.e., a compact patch of spins oriented against H may be stable, even if the energy could be lowered by flipping all them, because the flip of a single spin would rise the energy (actually, this flip may happen but is statistically reabsorbed in short times). The fate of the patch is governed by boundaries. A spin on a boundary of a patch feels a weaker local field, since some of its neighbors are oriented in the opposite. Straight boundaries in two or more dimensions separate spins that “know” the phase they belong to: since most of their neighbors are in that phase, the spins on the edges may flip more freely. Stripes that span the whole lattice are rather stable objects and may *resist* an opposite external field since spins that occasionally flip are surrounded by spins belonging to the opposite phase and therefore feel a strong local field that pushes them towards the phase opposed to the external field.

In one dimension with finite-range coupling, a single spin flip is able to create a “stripe” (perpendicularly to the lattice dimension) and therefore can destabilize the ordered phase. This

is the main reason for the absence of phase transitions in one dimension, unless the coupling extends on very large distances or some coupling is infinite (see the part on directed percolation, section “[Probabilistic Cellular Automata](#)”).

This model was proposed in the early 1920s by Lenz to Ising for his Ph.D. dissertation as a simple model of a ferromagnet. Ising studied it in one dimension, found that it shows no phase transition, and concluded (erroneously) that the same happened in higher dimensions. Most of the contemporaries rejected the model since it was not based on Heisenberg’s quantum mechanical model of ferromagnetic interactions. It was only in the forties that it started gaining popularity as a model of cooperative phenomena, a prototype of order-disorder transitions. Finally, in 1940, Onsager (1944) provided the exact solution of the two-dimensional Lenz-Ising model in zero external field. It was the first (and for many years the only) model exhibiting a nontrivial second-order transition whose behavior could be exactly computed.

Second-order transition has interested physicists for almost all the past century. In the vicinity of such transitions, the elements (say, spins) of the system are correlated up to very large distances. For instance, in the Lenz-Ising model (with coupling and more than one dimension), the high-temperature phase is disordered, and the low-temperature phase is almost completely ordered. In both these phases, the connected two-point correlation function

$$G_c(r) = \langle x_i x_{i+r} \rangle - \langle x_i \rangle^2$$

decreases exponentially, $G_c(r) \simeq \exp(-r/\xi)$, with $r = |\mathbf{r}|$. The length ξ is a measure of the typical size of patch of spins pointing in the same direction.

Near the critical temperature T_c , the correlation length ξ diverges like $\xi(T - T_c) \simeq (T - T_c)^{-\nu}$ ($\nu = 1$ for $d = 2$ and $\nu \simeq 0.627$ for $d = 3$, where d is the dimensionality of the system). In such case the correlation function is described by a power law

$$G_c(r) \simeq r^{2-d-\eta}$$

with $\eta = 1/4$ for $d = 1$ and $\eta \simeq 0.024$ for $d = 3$. This phase transition is an example of a *critical*

phenomenon (Binney et al. 1993); ν and η are examples of *critical exponents*.

The divergence of the correlation length indicates that there is no characteristic scale (ξ), and therefore fluctuations of all sizes appear. In this case, the details of the interactions are not so important, so that many different models behave in the same way, for what concerns, for instance, the critical exponents. Therefore, models can be grouped into *universality classes*, whose details are essentially given by “robust” characteristics like the dimensionality of space and of the order parameter, the symmetries, etc.

The power-law behavior of the correlation function also indicates that if we perform a *rescaling* of the system, it would appear the same or, conversely, that one is unable to estimate the *distance* of a pattern by comparing the “typical size” of particulars. This scale invariance is typical of many natural phenomena, from clouds (whose height and size are hard to be estimated), to trees and other plant elements, lungs, brain, etc.

Many examples of power laws and collective behavior can be found in natural sciences (Sornette 2006). Differently from what happens in the Lenz-Ising model, in these cases there is no parameter (like the temperature) that has to be fine-tuned, so one speaks of *self-organized criticality* (Bak et al. 1987).

Since the Lenz-Ising model is so simple, exhibits a critical phase, and can be exactly solved (in some cases), it has become the playground for a variety of modifications and applications to various fields. Clearly, most of the modifications do not allow analytical treatment and have to be investigated numerically. The Monte Carlo method allows to add a temporal dimension to a statistical model (Kawasaki 1972), i.e., to transform stochastic integrals into averages over fictitious trajectories. Needless to say, the Lenz-Ising model is the standard test for every Monte Carlo beginner, and most of the techniques for accelerating the convergence of averages have been developed with this model in mind (Swendsen and Wang 1987).

Near a second-order phase transition, a physical system exhibits *critical slowing down*, i.e., it reacts to external perturbations with an extremely

slow dynamics, with a convergence time that increases with the system size. One can extend the definition of the correlation function including the time dimension: in the critical phase also the temporal correlation length diverges (as a power law). This happens also for the Lenz-Ising model using the Monte Carlo dynamics, unless very special techniques are used (Swendsen and Wang 1987). Therefore, the dynamical version of the Lenz-Ising model can be used also to investigate relaxational dynamics and how this is influenced by the characteristics of the energy landscape. In particular, if the coupling J_{ij} changes sign randomly for every couple of sites (or the field H has random sign for each site), the energy landscape becomes extremely rugged. When spins flip in order to align to the local field, they may invert the field felt by neighboring ones. This *frustration* effect is believed to be the basic mechanism of the large viscosity and memory effects of *glassy* substances (Dotsenko 1994; Mezard et al. 1987).

The rough energy landscape of glassy systems is also challenging for optimization methods, like *simulated annealing* (Kirkpatrick et al. 1983) and its “improved” cousin, *simulated tempering* (Marinari and Parisi 1992). Again, the Lenz-Ising model is the natural playground for these algorithms.

The dynamical Lenz-Ising model can be formulated such that each spin is updated in parallel (Barkema and MacFarland 1994) (with the precaution of dividing cells into sublattices, in order to keep the neighborhood of each cell fixed during updates). In this way, it is essentially a probabilistic cellular automata, as illustrated in section “[Probabilistic Cellular Automata](#).”

Cellular Automata

In the same period in which traditional computation was developed, in the early 1950s, John Von Neumann was interested in the logic basis of life and in particular in self-reproduction, and since the analysis of a self-reproduction automata following the rules of real physics was too difficult, he designed a playground (a cellular automaton) with just enough “physical rules” in order to make its analysis possible (von Neumann and Burks,

1966). It was however just a theoretical exercise; the automaton was so huge that up to now it has not yet completely implemented (Von Neumann universal constructor 2008).

The idea of cellular automata is quite simple: take a lattice (or a graph) and put on each cell an automaton (all automata are equal). Each automaton exhibits its “state” (which is one out of a small number) and is programmed to react (change state) according to the state of neighbors and its present one (the *evolution rule*). All automata update their state synchronously.

Cellular automata share many similarities with the parallel version of the Lenz-Ising model. Differently from that, their dynamics is not derived from an energy, but is defined in terms of the transition rules. These rules may be deterministic or probabilistic. In the first case (illustrated in this section), cellular automata are *fully discrete, extended dynamical systems*. Probabilistic cellular automata are illustrated in section “[Probabilistic Cellular Automata](#).” The temporal evolution of deterministic cellular automata can be computed *exactly* (regardless of any approximation) on a standard computer.

Let us illustrate the simplest case, *elementary cellular automata*, in Wolfram’s jargon (1983). The lattice here is one-dimensional, so to identify an automaton it is sufficient to give one coordinate, say i , with $i = 1, \dots, N$. The state of the automaton on cell i at time t is represented by a single variable, $x_i(t)$, that can take only two values, “dead/live” or “inactive/active” or 0/1. The time is also discrete, so $t = 1, 2, \dots$

The parallel evolution of each automaton is given by the rule

$$x_i(t+1) = f(x_{i-1}(t), x_i(t), x_{i+1}(t)).$$

Since $x_i = 0, 1$, there are only eight possible combinations of the triple $\{x_{i-1}(t), x_i(t), x_{i+1}(t)\}$, from $\{0, 0, 0\}$ to $\{1, 1, 1\}$. For each of them, $f(x_{i-1}(t), x_i(t), x_{i+1}(t))$ is either zero or one, so the function f can be simply coded as a vector of eight bits, each position labeled by a different configuration of inputs. Therefore, there are only $2^8 = 256$ different elementary cellular automata that have been studied carefully (see, for instance, Wolfram (1983)).

In spite of their simplicity, elementary cellular automata exhibit a large variety of behaviors. In the language of dynamical systems, they can be “visually” classified (Wolfram 1983) as fixed points (class 1), limit cycles (class 2), and “chaotic” oscillations (class 3). A fourth class, “complex” CA, exhibits areas of repetitive or stable configurations with structures that interact with each other in complicated ways. A classic example is the *Game of Life* (Berlekamp et al. 1982). This two-dimensional cellular automaton is based on a simple rule. A cell may be either *dead* (0) or *alive* (1). A living cell survives if, among its 8 nearest neighbors, there are two or three alive cells, otherwise it dies and disappears. Generation is implemented through a rule for empty cells: they may become alive if surrounded by exactly three living cells. In spite of the simplicity of the rule, this automaton generates complex and long-living patterns, some of them illustrated in Fig. 2.


Complex CA have large transients, during which interesting structures may emerge. They finally relax into class-1 automata. It has been conjectured that they are able of computation, i.e., that one can “design” a universal computer using these CA as building blocks, as has been proved to be possible with the Game of Life. Another hypothesis, again confirmed by the

Game of Life, is that these automata are “near the edge” of self-organizing complexity. One can slightly “randomize” the Game of Life, allowing sometimes an exception to the rule. Let us introduce a parameter p that measures this randomness, with the assumption that $p = 0$ is the true “life.” Well, it was shown (Nordfalk and Alstrøm 1996) that the resulting model exhibits a second-order phase transition for a value of p very near zero.

Deterministic cellular automata have been investigated as prototypes of discrete dynamical systems, in particular for what concerns the definition of chaos. Visually, one is tempted to use this word also to denote the irregular behavior of “class-3” rules. However, the usual definition of chaos involves the sensitivity to an infinitesimally small perturbation: following the time dynamics of two initially close configurations, one can observe an amplification of their distance. If the initial distance (δ_0) is infinitesimal, then the distance grows exponentially for some time ($\delta(t) \simeq \delta_0 \exp(\lambda t)$), after which it tends to saturate, since the trajectories are generally bounded inside an attractor or due to the dimensions of the accessible space. The exponent λ depends on the initial configuration, and if this behavior is observed for different portions of the trajectory, it fluctuates: a trajectory spends some time in

Interaction-Based Computing in Physics,

Fig. 2 Some of the most common “animals” in the game of life, with the probability of encountering them in an asymptotic configuration (Bagnoli et al. 1991)



regions of high chaoticity, after which may pass through “quiet” zones. If one “renormalizes” periodically this distance, considering one system as the “master” and the other as a measuring device, one can accumulate good statistics and define a Lyapunov exponent λ that gives indications about the chaoticity of the trajectory, through a limiting procedure.

The accuracy of computation poses some problems. Since in computer numbers are always approximate, one cannot follow just “one” trajectory. The small approximations accumulate exponentially, and the computer time series actually *jumps* among neighboring trajectories. Since the Lyapunov exponent is generally not so sensible to a change of precision in computation, one can assume that the chaotic regions are rather compact and uniform, so that in general one associates a Lyapunov exponent to a system, not to an individual trajectory. Nevertheless, this definition cannot apply to completely discrete systems like cellular automata.

In any case, chaoticity is related to unpredictability. As first observed by Lorenz, and following the definition of Lyapunov exponent, the precision of an observation over a chaotic system is related to the average time for which predictions are possible. Like in weather forecasts, in order to increase the time span of a prediction, one has to increase the precision of the initial measurement. In extended system, this also implies to extend the measurements over a larger area. One can also consider a “synchronization” approach. Take two replicas of a system and let them evolve starting from different initial configurations. With a frequency and a strength that depends on a parameter q , one of these replicas is “pushed” towards the other one, so as to reduce their distance. Suppose that $q = 0$ is the case of no push and $q = 1$ is the case of extremal push, for which the two systems synchronize in a very short time. There should be a critical value q_c that separates these two behaviors (actually, the scenario may be more complex, with many phases (Bagnoli and Cecconi 2001)). In the vicinity of q_c , the distance between the two replicas is small, and the distance δ grows exponentially. The critical value q_c is such that the exponential growth compensates the shrinking factor and is therefore related to the Lyapunov exponent λ .

Finite-size cellular automata always follow periodic trajectories. Let us consider, for instance, Boolean automata of N cells. The number of possible different states is 2^N and due to determinism, once that a state has been visited twice the automata has entered a limit cycle (or a fixed point). One may have limit cycles with large basins of transient configurations (configurations that do not belong to the cycle). Many scenarios are possible. The set of different configurations may be divided in many basins, of small size (small transient) and small period, like in class-1 and class-2 automata. Or one may have large basins, with long transients that lead to short cycles, like in class-4 automata. Finally, one may have one or very few large basins, with long cycles that include most of the configurations belonging to the basin (small transients). This is the case of class-3 automata. For them, the typical period of a limit cycle grows exponentially, like the total number of configurations, with the system size, so that for moderately large system it is almost impossible to observe a whole cycle in a finite time. Another common characteristic of class-3 automata is that the configurations quickly decorrelate (in the sense of the correlation function) along a trajectory. If one takes into consideration as starting points two configurations that are the same but for a local difference, one observes that this difference amplifies and diffuses in class-3 automata, shrinks or remains limited in class 1 and class 2, and has an erratic transient behavior in class 4, followed by the fate of class 1 and class 2. Therefore, if one considers the possibility of not knowing exactly the initial configuration of an automaton, unpredictability grows with time also for such discrete systems. Actually, also (partially) continuous systems like coupled maps may exhibit this kind of behavior (Bagnoli and Cecconi 2001; Cecconi et al. 1998; Crutchfield and Kaneko 1988; Politi et al. 1993). Along this line, it is possible to define an equivalent of the Lyapunov exponents for CA (Bagnoli et al. 1992). The synchronization procedure can be applied also to cellular automata, and it correlates well with the Lyapunov exponents (Bagnoli and Rechtman 1999).

An “industrial” application of cellular automata is their use for modeling gases. The *hydrodynamical equations*, like the Navier-Stokes ones,

simply reflect the conservation of mass, momentum, and energy (i.e., rotational, translational, and time invariance) for the microscopic collision rules among particles. Since the modeling of a gas via molecular dynamics is rather cumbersome, some years ago it was proposed (Frisch et al. 1986; Hardy et al. 1973) to simplify drastically the microscopic dynamics using particles that may travel only along certain directions with some discrete velocities and jumping in discrete time only among nodes of a lattice, indeed, a cellular automaton. It has been shown that their macroscopic dynamics is described by usual hydrodynamics laws (with some odd features related to the underlining lattice and finiteness of velocities) (Rothman and Zaleski 2004; Wolf-Gladrow 2004).

The hope was that these Lattice Gas Cellular Automata (LGCA) could be simulated so efficiently in hardware to make possible the investigation of turbulence, or, in other words, that they could constitute the Ising model of hydrodynamics. While they are indeed useful to investigate certain properties of gases (for instance, chemical reactions (Lawniczak et al. 1991), or the relationship between chaoticity and equilibrium (Bagnoli and Rechtman 2009)), they resulted too noisy and too viscous to be useful for the investigation of turbulence. Viscosity is related to the transport of momentum in a direction perpendicular to the momentum itself. If the collision rule does not “spread” quickly the particles, the viscosity is large. In LGCA there are many limitations to collisions, so that in order to lower viscosity one has to consider averages over large patches, thus lowering the efficiency of the method.

However, LGCA inspired a very interesting approximation. Let us consider a large assembly of replicas of the same system, each one starting from a different initial configuration, all compatible with the same macroscopic initial conditions. The macroscopic behavior after a certain time would be the *average* over the status of all these replicas. If one assumes a form of local equilibrium, i.e., applies the mean-field approximation for a given site, one may try to obtain the dynamics of the *average* distribution of particles, which in principle is the same of “exchanging” particles that happen to stay on the same node among replicas.

It is possible to express the dynamics of the average distribution in a simple form: it is the Lattice Boltzmann Equation (LBE) (Chopard et al. 2002; Succi 2001; Wolf-Gladrow 2004). The method retains many properties of LGCA like the possibility of considering irregular and varying boundaries and may be simulated in a very efficient way with parallel machines (Succi 2001). Differently from LGCA, there are numerical stability problems to be overcome.

Probabilistic Cellular Automata

In deterministic automata, given a local configuration, the future state of a cell is univocally determined. However, let us consider the case of measuring experimentally some pattern and trying to analyze it in terms of cellular automata. In time-series analysis, it is common to perform averages over spatial patches and temporal intervals and to discretize the resulting value. For instance, this is the natural result of using a camera to record the temporal evolution of an extended system, for instance the turbulent and laminar regions of a fluid. The resulting pattern symbolically represents the dynamics of the original system, and if it is possible to extract a “rule” out of this pattern, it would be extremely interesting for the construction of a model. In general, however, one observes that sometimes a local configuration is followed by a symbol, and sometimes the same local configuration is followed by another one. One should conclude that the neighborhood (the local configuration) does not univocally determine the following symbol.

One can extend the “range” of the rule, adding more neighbors farther in space and time (Rabiner 1989). By doing so, the “conflicts” generally reduce but at the price of increasing the complexity of the rule. At the extremum, one could have an automaton with infinite “memory” in time and space that perfectly reproduces the observed patterns but with almost nonpredictive power, since it is extremely unlucky that the same *huge* local configuration is encountered again.

So, one may prefer to limit the neighborhood to some finite extension and accept that the rule sometimes “outputs” a symbol and sometimes another one. One defines a local transition

probability $\tau(x_i(t+1)|X_i(t))$ of obtaining a certain symbol x_i at time $t+1$ given a local configuration X_i at time t . Deterministic cellular automata correspond to the case $\tau = 0, 1$. The parallel version of the Lenz-Ising model can be reinterpreted as a probabilistic cellular automaton.

Starting from the local transition probabilities, one can build up the transition probability $T(x|y)$ of obtaining a configuration x given a configuration y . $T(x|y)$ is given by the product of the local transition probabilities τ (Bagnoli 2000). One can read the configurations x and y as indexes, so that T can be considered as a matrix. The normalization of probability corresponds to the constraint $\sum_x T(x|y) = 1 \forall y$.

Denoting with $P(x, t)$ the probability of observing a given configuration x at time t , and with $P(t)$ the whole distribution at time t , we have for the evolution of the distribution

$$P(x, t+1) = T P(t),$$

with the usual rules for the product of matrices and vectors. Therefore, the transition matrix T defines a Markov process, and the asymptotic state of the system is given by the eigenvalues of T . The largest eigenvalue is always 1, due to the normalization of the probability distribution, and the corresponding eigenvector is the asymptotic distribution. The theory of Markov chains says that if T is finite and irreducible, i.e., it cannot be rewritten (renumbering rows and columns) as blocks of noninteracting subspaces, then the second eigenvalue is strictly less than one and the asymptotic state is unique. In this case, the second eigenvalue determines the convergence time to the asymptotic state. For finite systems, often the matrix T is irreducible. However, in the limit of infinite size, the largest eigenvalue may become degenerate, and therefore there are more than one asymptotic state. This is the equivalent of a phase transition for Markov processes.

For the parallel Lenz-Ising model, the elements of the matrix T are given by the product of local transition rules of the Monte Carlo dynamics. They depend on the choice of the algorithm but essentially have the form of exponentials of the difference in energy, divided by the temperature.


Although a definitive proof is still missing, it is plausible that matrices with all elements different from zero correspond to some equilibrium model, whose transition rules can be derived from an energy function (Georges and le Doussal 1989).

Since probabilistic cellular automata are defined in terms of the transition probabilities, one is free to investigate models that go beyond equilibrium. For instance, if some transition probability takes the value zero or one, in the language of equilibrium system, this would correspond to some coupling (like the coupling factor J of the Lenz-Ising model) that becomes infinite. This case is not so uncommon in modeling. The inverse of a transition probability corresponds to the average waiting time for the transition to occur in a continuous-time model (one may think to chemical reactions). Some transitions may have a waiting time so long with respect to the observation interval, to be practically irreversible. Therefore, probabilistic cellular automata (alongside other approaches like for instance annihilating random walks) allow the exploration of out-of-equilibrium phenomena.

One such phenomenon is directed percolation, i.e., a percolation process with a special direction (time) along which links can only be crossed one way (Broadbent and Hammersley 1957). Let us think, for instance, of the spreading of an infection in a one-dimensional lattice, with immediate recovery (SIS model). An ill individual can infect one or both of his two neighbors but returns to the susceptible state after one step. The paths of infection (see Fig. 3) can wander in the space directions but are directed in the time directions.

The parallel version of a directed percolation process can be mapped onto probabilistic cellular automata. The simplest case, in one spatial dimension and with just two neighbors, is called the Domany-Kinzel model (Domany and Kinzel 1984). It is even more general than the usual directed percolation, allowing “nonlinear” interactions among sites in the neighborhood (e.g., two wet sites may have less probability of percolating than one alone).

These processes are interesting because there is an *absorbing state* (Bagnoli et al. 2001; Hinrichsen 1997), which is the dry state for the wetting phenomenon and the healthy state for the spreading of



Interaction-Based Computing in Physics, Fig. 3 An example of a directed percolation cluster

epidemics. Once the system has entered this absorbing state, it cannot exit, since the spontaneous appearing of a wet site or of an ill individual is forbidden. For finite systems, the theory of Markov chains says that the system will “encounter,” sooner or later, this state that therefore corresponds to the unique asymptotic state. For infinitely extended systems, a phase transition can occur, for which wet sites percolate for all “times,” the epidemics become endemic, and so on.

Again these are examples of *critical phenomena*, with exponents different from the equilibrium case.

Agent-Based Simulations

Cellular automata can be useful in modeling phenomena that can be described in lattice terms. However, many phenomena require *moving particles*. Examples may be chemical reactions, ecological simulations, and social models. When particles are required to obey hydrodynamics constraints, i.e., to collide conserving mass, momentum, and energy, one can use lattice gas cellular automata or the lattice Boltzmann equation. However, in general one is interested in modeling just a macroscopic scale, assuming that what happens at lower level is just “noise”. In accordance with the degree of complexity (or “intelligence”) assigned to particles, one can develop models based on the

concept of walkers that move more or less randomly. From the simulation point of view, walkers are not very different from the graphs succinctly described above. In this case, the identifier i is just a label that allows to access walker’s data, among which there are the coordinates of the walker (that may be continuous), its status, and so on.

In order to let walkers interact, one is interested in finding efficiently all walkers that are nearer than a given distance from the one under investigation. This is the same problem one is faced with when developing codes for molecular dynamics simulations (Rapaport 2004): scanning all walkers in order to compute their mutual distance is a procedure that grows quadratically with the number of walkers. One “trick” is that of dividing the space in cells, i.e., to define an associated lattice. Each cell contains a list of all walkers that are located inside it. In this way, one can directly access all walkers that are in the same or neighboring cells of the one under investigation.

Moreover, one can exploit the presence of the lattice to implement on it a “cellular automaton: that may interact with walkers, in order to simulate the evolution of “fields”. Just for example, the simulation of a herd of herbivores that move according to the exhaustion of the grass and the parallel growing of the vegetation can be modeled associating the “grass” to a cellular automaton and the herbivores to walkers. This simulation scheme is quite flexible, allowing to implement random and deterministic displacements of moving objects or agents, continuous or discrete evolution of “cellular objects,” and their interactions. Many simulation tools and games are based on this scheme (Repat – recursive porous agent simulation toolkit 2008; Wilensky 1999), and they generally allow the contemporary visualization of a graphical representation of the state of the system. They are valuable didactic tool and may be used to “experiment” with these artificial worlds. As often the case, the flexibility is paid in terms of efficiency and speed of simulation.

Future Directions

Complex systems, like, for instance, human societies, cannot be simulated starting from *physical*

laws. This is sometimes considered a weakness of the whole idea of studying such systems from a quantitative point of view. We have tried to show that actually even the “hardest” discipline, physics, always deals with models that have finally to be simulated on computers, making various assumptions and approximations. Theoretical physics is accustomed for a long time to *extreme simplifications* of models, hoping to enucleate the *fundamental ingredients* of a complex behavior. This approach has proved to be quite rewarding for our understanding of nature.

In recent years, physics have been seen studying many fields not traditionally associated to physics: molecular biology, ecology, evolution theory, neurosciences, psychology, sociology, linguistics, and so on. Actually, the word “physics” may refer either to the classical subjects of study (mainly atomic and subatomic phenomena, structure of matter, cosmological, and astronomical topics) or to the “spirit” of the investigation that may apply to almost any discipline. This spirit is essentially that of building simplified quantitative models, composed by many elements, and study them with theoretical instruments (most of the times, applying some form of mean-field treatment) and with computer simulations.

This approach has been fruitful in chemistry and molecular biology, and nowadays many physical journals have sections devoted to *multi-disciplinary studies*. The interesting thing is that not only physicists have brought some *mathematics* into fields that are traditionally more *qualitative* (which often corresponds to *linear* modeling, plus noise), but have also discovered many interesting questions to be investigated and new models to be studied. One example is given by the current investigations about the structure of social networks that were “discovered” by physicists in the nontraditional field of social studies.

Another contribution of physicists to this “new way” of performing investigations is the use of networked computers. For a long time, physicists have used computers in performing computations, storing data and diffusing information using the Internet. Actually, the concept of what is now the World Wide Web was born at CERN, as a method for sharing information among laboratories

(Cailliau 1995). The high-energy physics experiments require a lot of simulations and data processing, and physicists (among others) developed protocols to distribute this load on a grid of networked computers. Nowadays, a European project aims to “open” grid computing to other sciences (European Grid Infrastructure).

It is expected that this combination of quantitative modeling and grid computing will stimulate innovative studies in many fields. Here is a small sparse list of possible candidates:

- Theory of evolution, especially for what concerns evolutionary medicine
- Social epidemiology, coevolution of diseases and human populations, interplay between sociology and epidemics
- Molecular biology and drug design, again driven by medical applications
- Psychology and neural sciences: it is expected that the “black box” of traditional psychology and psychiatry will be replaced by explicit models based on brain studies
- Industrial and material design
- Earth sciences, especially meteorology, volcanology, and seismology
- Archaeology: simulation of ancient societies, reconstruction of historical and prehistorical climates

Nowadays, the term *cellular automata* has enlarged its meaning, including any system whose elements do not move (in opposition to *agent-based modeling*). Therefore, we now have cellular automata on nonregular lattices, non-homogeneous, with probabilistic dynamics (see section “Probabilistic Cellular Automata”), etc. (El Yacouby et al. 2006). They are therefore considered more as a “philosophy” of modeling rather than a single tool. In some sense, cellular automata (and agent-based) modeling is opposed to the spirit of describing a phenomena using differential equations (or partial differential equations). One of the reasons is that the language of *automata* and *agents* is simpler and requires less training than that of differential equations. Another reason is that at the very end, any *reasonable* problem has to be investigated using

computers, and while the implementation using discrete elements is straightforward (even if careful planning may speed up dramatically the simulation), the computation of partially differential equations is an art in itself.

However, the final success of this approach is related to the availability of high-quality experimental data that allow to discriminate among the almost infinite number of models that can be built.

Bibliography

Primary Literature

- Albert R, Barabasi AL (2002) Statistical mechanics of complex networks. *Rev Mod Phys* 74:47–97
- Bagnoli F, Rechtman R, Ruffo S (1991) Some facts of life. *Physica A* 171:249–264
- Bagnoli F (2000) Cellular automata. In: Bagnoli F, Ruffo S (eds) *Dynamical modeling in biotechnologies*. World Scientific, Singapore, p 1
- Bagnoli F, Cecconi F (2001) Synchronization of non-chaotic dynamical systems. *Phys Lett A* 282(1–2):9–17
- Bagnoli F, Rechtman R (1999) Synchronization and maximum Lyapunov exponents of cellular automata. *Phys Rev E* 59(2):R1307–R1310
- Bagnoli F, Rechtman R (2009) Thermodynamic entropy and chaos in a discrete hydrodynamical system *Phys Rev E* 79:041115
- Bagnoli F, Rechtman R, Ruffo S (1992) Damage spreading and Lyapunov exponents in cellular automata. *Phys Lett A* 172:34
- Bagnoli F, Boccara N, Rechtman R (2001) Nature of phase transitions in a probabilistic cellular automaton with two absorbing states. *Phys Rev E* 63(4):046116
- Bak P, Tang C, Wiesenfeld K (1987) Self-organizing criticality: an explanation of $1/f$ noise. *Phys Rev A* 38:364–374
- Barkema GT, MacFarland T (1994) Parallel simulation of the Ising model. *Phys Rev E* 50(2):1623–1628
- Berlekamp E, Conway J, Guy R (1982) *What is life? Games in particular*, vol 2. Academic, London. Chap. 25
- Binney J, Dowrick N, Fisher A, Newman MEJ (1993) *The theory of critical phenomena*. Oxford Science/Clarendon Press, Oxford
- Boccaletti S, Latora V, Moreno Y, Chavez M, Hwang DU (2006) Complex networks: structure and dynamics. *Phys Rep* 424(4–5):175–308
- Broadbent S, Hammersley J (1957) Percolation processes I. Crystals and mazes. *Proc Camb Philos Soc* 53:629–641
- Cailliau R (1995) A short history of the web. http://www.netvalley.com/archives/mirrors/robert_cailliau_speech.htm. Accessed 10 Apr 2017
- Car R, Parrinello M (1985) Unified approach for molecular dynamics and density-functional theory. *Phys Rev Lett* 55(22):2471–2474
- Cecconi F, Livi R, Politi A (1998) Fuzzy transition region in a one-dimensional coupled-stable-map lattice. *Phys Rev E* 57(3):2703–2712
- Chopard B, Luthi P, Masselot A, Dupuis A (2002) Cellular automata and lattice Boltzmann techniques: an approach to model and simulate complex systems. *Adv Complex Syst* 5(2):103–246
- Crutchfield J, Kaneko K (1988) Are attractors relevant to turbulence? *Phys Rev Lett* 60(26):2715–2718
- Daxois T, Peyrard M, Ruffo S (2005) The Fermi-Pasta-Ulam ‘numerical experiment’: history and pedagogical perspectives. *Eur J Phys* 26:S3–S11
- Domany E, Kinzel W (1984) Equivalence of cellular automata to Ising models and directed percolation. *Phys Rev Lett* 53(4):311–314
- Dotsenko V (1994) *An introduction to the theory of spin glasses and neural networks*. World Scientific, Singapore
- El Yacouby S, Chopard B, Bandini S (eds) (2006) *Cellular automata, Lecture notes in computer science*, vol 4173. Springer, Berlin
- European Grid Infrastructure. <https://www.egi.eu/>. Accessed 10 Apr 2017
- Fermi E, Pasta J, Ulam S (1955) Los Alamos report la-1940. In: Segré E (ed) *Collected papers of Enrico Fermi*. University of Chicago Press, Chicago
- Frisch U, Hasslacher B, Pomeau Y (1986) Lattice-gas automata for the Navier-Stokes equation. *Phys Rev Lett* 56(14):1505–1508
- Gardiner CW (1994) *Handbook of stochastic methods for physics, chemistry, and the natural sciences*, Springer series in synergetics, vol 13. Springer, Berlin
- Georges A, le Doussal P (1989) From equilibrium spin models to probabilistic cellular automata. *J Stat Phys* 54(3–4):1011–1064
- Hardy J, Pomeau Y, de Pazzis O (1973) Time evolution of a two-dimensional classical lattice system. *Phys Rev Lett* 31(5):276–279
- Harlow H, Metropolis N (1983) Computing & computers – weapons simulation leads to the computer era. *Los Alamos Sci* 4(7):132
- Haw M (2005) Einstein’s random walk. *Phys World* 18:19–22
- Hinrichsen H (1997) Stochastic lattice models with several absorbing states. *Phys Rev E* 55(1):219–226
- Jaynes E (1957) Information theory and statistical mechanics. *Phys Rev* 106(4):620–630
- Kaneko K (1985) Spatiotemporal intermittency in coupled map lattices. *Progr Theor Phys* 74(5):1033–1044
- Kawasaki K (1972) Kinetics of Ising model. In: Domb CM, Green MS (eds) *Phase transitions and critical phenomena*, vol 2. Academic, New York, p 443
- Kirkpatrick S, Gelatt CG Jr, Vecchi MP (1983) Optimization by simulated annealing. *Science* 220:671–680
- Lawniczak A, Dab D, Kapral R, Boon JP (1991) Reactive lattice gas automata. *Phys D* 47(1–2):132–158
- Marinari E, Parisi G (1992) Simulated tempering: a new Monte Carlo scheme. *Europhys Lett* 19:451–458
- May R (1976) Simple mathematical models with very complicated dynamics. *Nature* 261:459–467

- Metropolis N, Hewlett J, Rota GC (eds) (1980) A history of computing in the twentieth century. Academic, New York
- Mezard M, Parisi G, Virasoro MA (1987) Spin glass theory and beyond. World scientific lecture notes in physics, vol 9. World Scientific, Singapore
- Newman ME (2005) Power laws, Pareto distributions and Zipf's law. *Contemp Phys* 46:323–351
- Niss M (2005) History of the Lenz-Ising model 1920–1950: from ferromagnetic to cooperative phenomena. *Arch Hist Exact Sci* 59(3):267–318
- Nordfalk J, Alstrøm P (1996) Phase transitions near the “game of life”. *Phys Rev E* 54(2):R1025–R1028
- Onsager L (1944) Crystal statistics. I. A two-dimensional model with an order-disorder transition. *Phys Rev* 65:117–149
- Oestreicher C (2007) A history of chaos theory. *Dialogues Clin Neurosci* 9(3):279–289. Available online <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3202497/>
- Politi A, Livi R, Oppo GL, Kapral R (1993) Unpredictable behaviour of stable systems. *Europhys Lett* 22(8):571–576
- Rabiner L (1989) A tutorial on hidden Markov models and selected applications in speech recognition. *Proc IEEE* 77(2):257–286
- Rapaport DC (2004) The art of molecular dynamics simulation. Cambridge University Press, Cambridge
- Repast – recursive porous agent simulation toolkit (2008) <http://repast.sourceforge.net/>. Accessed 10 Apr 2017
- Rothman DH, Zaleski S (2004) Lattice-gas cellular automata. Monographs and texts in statistical physics. Collection Alea-Saclay, Paris
- Sornette D (2006) Critical phenomena in natural sciences, Springer series in synergetics. Springer, Berlin
- Stauffer D, Aharony A (1994) Introduction to percolation theory. Taylor Francis, London
- Succi S (2001) The lattice Boltzmann equation for fluid dynamics and beyond. Numerical mathematics and scientific computation. Oxford University Press, Oxford
- Swendsen R, Wang JS (1987) Nonuniversal critical dynamics in Monte Carlo simulations. *Phys Rev Lett* 58(2):86–88
- van Kampen NG (1992) Stochastic processes in physics and chemistry. North-Holland, Amsterdam
- von Neumann J, Burks AW (1966) Theory of self-reproducing automata. University of Illinois Press, Urbana/London
- Von Neumann universal constructor (2008) http://en.wikipedia.org/wiki/Von_Neumann_Universal_Constructor. Accessed 10 Apr 2017
- Watts D, Strogatz SH (1998) Collective dynamics of ‘small-world’ networks. *Nature* 393:440–441
- Wilensky U (1999) Netlogo. Center for connected learning and computer-based modeling, Northwestern University, Evanston. <http://ccl.northwestern.edu/netlogo/>. Accessed 10 Apr 2017
- Wolf-Gladrow D (2004) Lattice-gas cellular automata and lattice Boltzmann models: an introduction, Lecture notes in mathematics, vol 1725. Springer, Berlin
- Wolfram S (1983) Statistical mechanics of cellular automata. *Rev Mod Phys* 55:601–644

Books and Reviews

- Boccara N (2004) Modeling complex systems. In: Graduate texts in contemporary physics. Springer, Berlin
- Bungartz H-J, Mundani R-P, Frank AC (2005) Bubbles, jaws, moose tests, and more: the wonderful world of numerical simulation, Springer VideoMATH. Springer, Berlin. (DVD)
- Chopard B, Droz M (2005) Cellular automata modeling of physical systems. In: Collection Alea-Saclay: monographs and texts in statistical physics. Cambridge University Press, Cambridge
- Deisboeck S, Kresh JY (2006) Complex systems science in biomedicine. In: Deisboeck S, Kresh JY (eds) Topics in biomedical engineering. Springer, New York
- Gould H, Tobochnik J, Christian W (2007) An introduction to computer simulation methods: applications to physical systems. Addison-Wesley, New York
- Landau RH (2005) A first course in scientific computing: symbolic, graphic, and numeric modeling using maple, java, Mathematica, and Fortran90. Princeton University Press, Princeton
- Open Source Physics. <http://www.opensourcephysics.org/>. Accessed 10 Apr 2017
- Resnick M (1994) Turtles, termites, and traffic jams. Explorations in massively parallel microworlds. In: Complex adaptive systems. MIT Press, Cambridge
- Shalizi C Cosma's home page. <http://www.cscs.umich.edu/~crshalizi/>. Accessed 10 Apr 2017

Swarm Intelligence

Gerardo Beni

University of California, Riverside, CA, USA

Article Outline

Glossary

Definition of the Subject and Its Importance

Introduction

Biological Systems

Robotic Systems

Artificial Life Systems

Definition of Swarm

Standard-Mathematics Methods

Swarm Optimization

Particle Swarm Optimization (PSO)

Ant Colony Optimization (ACO)

Nonlinear Differential Equation Methods

Limitations of Standard-Mathematics Methods

Cellular-Computing Methods

Intelligence as Universal Computation

Relations to Standard-Mathematics Methods

Randomness in Swarm Intelligence

The Implicit Assumption of Asynchrony
Irrelevance

Asynchronous Swarms

Types of Asynchrony

Modeling Asynchrony by Synchronous Swarms

Local Synchrony and Self-Synchronization

The Natural Asynchrony of Swarms

The Realization of Asynchronous Swarms

Characteristics of Swarm Intelligence

Dynamics in Swarm Intelligence

Unpredictability in Swarm Intelligence

Swarms of Intelligent Units

Future Directions

Bibliography

Glossary

Ant colony optimization Probabilistic optimization algorithm where a colony of artificial ants cooperate in finding solutions to optimization problems.

Cellular automaton A system evolving in discrete time steps, with four properties: a grid of cells, a set of possible states of the cells, a neighborhood, and a function which assigns a new state to a cell given the state of the cell and of its neighborhood.

Cellular-computing architecture Computer design that uses cellular automata and related machines, as processors and as storage of instruction and data.

Dynamic cellular-computing system Cellular-computing system whose cells are mobile. Elementary swarmAn ordered set of N units described by the N components v_i ($i = 1, 2, \dots, N$) of a vector v ; any unit i may update the vector at any time t_i , using a function f of K_i vector components. $\forall i \in N: v_i(t+1) = f(v_{k \in K(i)}(t))$.

Game of life A cellular automaton designed to simulate lifelike phenomena.

Intelligence (working definition for swarm intelligence) Ability to carry out universal computation.

Natural asynchrony Asynchronous updating characterized by three properties: more than one unit may update at each time step, any unit may update more than once in each updating cycle, and the updating order varies randomly for every updating cycle.

Optimization algorithms Algorithms to satisfy a set of constraints and/or optimize (e.g., minimize) a function by systematically choosing the values of the variables from an allowed set.

Particle swarm optimization Probabilistic optimization algorithm where a swarm of potential

solutions (particles) cooperate in finding solutions to discrete optimization problems.

Pheromone A chemical that triggers an innate behavioral response in another member of the same animal species.

Stigmergy Indirect communication through modification of the environment.

Swarm intelligence Definition 1 (section “[Definition of the Subject and Its Importance](#)”): The *intuitive* notion of “swarm intelligence” is that of a “swarm” of agents (*biological or artificial*) which, without central control, collectively (and only collectively) carry out (unknowingly, and in a somewhat-random way) tasks normally requiring some form of “intelligence.” Definition 5 (section “[Swarms of Intelligent Units](#)”): The capability of universal computation carried out with natural asynchrony by a dynamic cellular-computing system, none of whose cells can predict the computation done by the swarm.

Swarm optimization Ant colony optimization particle swarm optimization and related probabilistic optimization algorithms.

Swarm robotics The technology of robotic systems capable of swarm intelligence.

Unpredictable system A system such that complete knowledge of its state and operation at any given time is insufficient to compute the system’s future state before the system reaches it.

Von Neumann architecture Computer design that uses one processing unit and one storage unit holding both instructions and data.

Definition of the Subject and Its Importance

The research area identified as “swarm intelligence” (SI) has been evolving now for 30 years. The term “swarm intelligence” first appeared in 1989 (Beni and Wang 1989a, b). By 2007 (when the first version of this entry appeared), “swarm intelligence” was in the title of four books (Bonabeau et al. 1999; Kennedy et al. 2001; Abraham et al. 2006; Engelbrecht 2006), in two series of conference proceedings (Dorigo et al.

2006; IEEE 2007), and in a new technical journal (2007), without mentioning other areas in which the term swarm itself had become popular. At the time of the first update (mid-2013), there were already dozens of books on or related to swarm intelligence. There are now (mid-2019) five journals dedicated strictly to swarm intelligence. A search for “Swarm Intelligence” on the Internet yields about 18 million results.

As the use of the term “swarm intelligence” has spread, its meaning has broadened to a point in which it is often understood to encompass almost any type of collective behavior. And since the term “swarm intelligence” has popular appeal, it is also sometimes used in contexts which have limited scientific or technological content. Some meanings, however, refer rigorously to precise concepts. The following treatment of swarm intelligence is based only on concepts that can be clearly defined and quantified. Hence, it is more restricted than some broader swarm intelligence presentations, but, even so, it describes an inter-related scientific/technical core which forms a solid basis for a well-defined multidisciplinary research area.

Definition 1 The *intuitive* notion of “swarm intelligence” is that of a “swarm” of agents (*biological or artificial*) which, without central control, collectively (and only collectively) carry out (unknowingly, and in a somewhat-random way) tasks normally requiring some form of “intelligence.”

A more specific definition requires a detailed discussion, and so it is given at the end of the article. (Sections “[Characteristics of Swarm Intelligence](#)” and “[Swarms of Intelligent Units](#)”.)

Although this notion of swarm intelligence might seem vague, we will see in the course of this entry that in fact it has many specific implications. Note that the notion is broad, which partly explains its widespread use, but not so broad as to include any type of collective action of groups of simple entities, as will become clear later.

These characteristics of swarm intelligence are also those of several biological systems, e.g., some insect societies or some components of the

immune system, so that swarm intelligence has become important for understanding certain mechanisms in biology.

Technologically, the importance of “swarms” is mainly based on the potential advantages over centralized systems. The potential advantages are:

1. *Economy*: the swarm components (units) are simple, hence (in principle) mass producible, modularizable, interchangeable, and disposable.
2. *Reliability*: due to the redundancy of the components, destruction/death of some units has negligible effect on the accomplishment of the task, as the swarm adapts to the loss of few units.
3. Ability to perform *tasks beyond those of centralized systems*, e.g., escaping enemy detection.

From this initial perspective on potential advantages, the actual application of swarm intelligence has extended to many areas, described in the body of this entry, and its potential for future applications remains high, as discussed in the concluding section.

Some current, proposed, and/or potential applications are in defense and space technologies (e.g., control of groups of unmanned vehicles in land, water, or air), flexible manufacturing systems, advanced computer technologies (biocomputing), medical technologies, and telecommunications.

SI can be seen also as part of the broader field of computational intelligence (CI) (Keller et al. 2016), which comprises computational methods for problems not solvable by the first principles or statistical modeling. Since such problems (often ill-posed inverse problems with a high level of noise and uncertainty) appear to have been solved by biological systems, CI attempts to reproduce, to some extent, these natural methods of solution (Kacprzyk and Pedrycz 2015). Swarm intelligence is one such method of CI along with more mature approaches such as artificial neural networks, fuzzy logic, and evolutionary computation. For an example of current CI studies, see Mandal and Devadutta (2019).

Introduction

Swarm intelligence (SI) investigations were initially motivated by studies of groups of simple robotic units which offered promise for technology. These “swarms” were modeled as collections of simple quasi-identical units with decentralized control and independent clocks (Beni and Wang 1989a). The number of units is intermediate between those of typical systems investigated in physics and other traditional science fields; in fact, the swarm is of the order of 10^2 to 10^8 units, where $s \ll 23$, i.e., the swarm is not composed of so many units that statistical physics methods can be applied to it nor of such a few units that its dynamic can be solved exactly (or numerically to high precision).

These features are typical also of many biological systems, such as insect societies. They are also the features of robotic systems potentially economic, reliable, and capable of tasks beyond the capabilities of centralized systems.

There are various established fields of science and technology that deal, to a certain extent, with the intuitive notion of SI: artificial life, ethology, robotics, artificial intelligence, computation, self-organization, complexity, economics, and sociology. In all these fields, there is at some level, or in some application, the need to understand, model, and predict the operation of groups of units which only by working together (in a not very structured way and “unaware” of the evolution of the group) carries out “intelligent” tasks. A simple illustration is provided by classic economics. In a free market economy, people trade with each other in an unstructured way: each makes independent decisions at unpredictable times and unaware of the global results. But the outcome is the solution to a complex problem – the problem of correct pricing. Such a “swarm” solves a problem not (or poorly) solvable by centralized control economies, thus exhibiting, in a certain sense, a high form of “intelligence.”

From this example, it is clear that the intuitive notion of SI is easy to grasp. But it is not so easy to make it less vague and more precise and quantitative, i.e., to make it a useful working concept for science and technology. That the concept of SI is

not easy to quantify follows from the difficulty of defining several of the key components of the intuitive notion of SI (Definition 1). First, “intelligence” is a notoriously ambiguous concept. Second, “randomly” is also not easily defined and quantified. Third, “only collectively” must be specified in terms of the critical number of agents required for the emergence of SI. In what sense is a unit “simple” or “unintelligent” and the task carried out “complex” or “intelligent”? Fourth, “unknowingly” implies that the global status and the goal of the swarm are, at least to some extent, unknown to the single agents. Which algorithms and communication schemes result in tasks carried out “unknowingly” by the agents?

Because of these difficulties, in this entry, we first use the aforementioned (Definition 1) intuitive notion of SI to describe the current main areas of studies considered to be SI. This will provide an overview of the current status of the field; it will make it possible to quantify the four vague concepts (“intelligence,” “randomly,” “collectively,” “unknowingly”) and to reach a more sharply defined concept of SI. From this, we will be able to see more clearly the limitations of SI and so its realistic potential for future applications.

In this entry, the main areas of SI studies are described by making three very broad distinctions: (1) scientific interest versus technological interest (sections “Biological Systems,” “Robotic Systems,” “Artificial Life Systems,” and “Definition of Swarm”); (2) standard mathematics versus cellular computational mathematics (sections “Standard-Mathematics Methods,” “Swarm Optimization,” “Nonlinear Differential Equation Methods,” “Limitations of Standard-Mathematics Methods,” “Cellular-Computing Methods”); and (3) synchronous operation versus asynchronous operation (sections “Randomness in Swarm Intelligence,” “Asynchronous Swarms,” “The Realization of Asynchronous Swarms,” “Characteristics of Swarm Intelligence,” “Swarms of Intelligent Units”). These distinctions in turn will provide a guide to clarifying the four vague concepts in the intuitive definition of SI (Definition 1) and, thus, a conceptual orientation for future studies and applications; they will also provide criteria for evaluating the promise of SI to solve complex problems that traditional approaches cannot.


Focusing on the first distinction (scientific vs. technological interest), the main scientific interest in SI originated with the work of biologists studying insect societies (Bonabeau et al. 1999). The main technological interest originated with roboticists trying to design distributed robotic systems (Beni and Wang 1989a, b). A valuable reference on the development of SI is Bonabeau et al. (1999), dealing, in parallel, with these two interrelated interests. For a more recent survey, see Tan (2018).

Biological Systems

Probably the best-known and seminal biology experiment in SI is the “double bridge” experiment by Goss et al. (1989). While studying the foraging of ants, they observed that if ants, starting from a point S, could reach food at a point F via two paths of different lengths, the ants would at first choose one of the two paths randomly; but after some ants had returned from F to S, more ants would choose to go from S to F via the shortest path; and eventually practically all the ants would choose the shortest path (see Fig. 1).

The ants following the shorter path (the lower path) return to the source before the ants which have taken the longer path. In this way, the shorter path has a higher density of pheromones; as a result, ants starting at S will now prefer the shorter path.

The key insight was the realization that the ants were finding the best path via *stigmergy*, that is, by *communicating through modification of the environment*. Ants are blind, but they communicate chemically via pheromones. By laying



Swarm Intelligence, Fig. 1 Illustration of the double-bridge experiments

pheromones along the path when returning from the food source F, the ants effectively marked the shortest path by laying more pheromones on it. After that, the ants that would start from S would choose the path marked by more pheromones, i.e., the shortest path.

Thus, a method of self-organization and a method to solve a nontrivial problem by a form of collective intelligence, with many of the elements of the intuitive definition of SI given above, were observed and understood.

Later, Dorigo (1992) realized that this method could be abstracted and generalized to design algorithms that rely on “artificial ants” to solve much more complex problems (see section “[Swarm Optimization](#)”). Thus, the close connection between biological studies of SI and its potential for technological application was first clearly demonstrated.

Many other experiments in ants and other social insects have confirmed the potential for developing bioinspired algorithms (Bonabeau et al. 1999; Dorigo and Stutzle 2004; Olariu and Zomaya 2005; Passino 2004). For actual insect societies, various ant algorithms have been applied to model tasks such as division of labor, cemetery organization, brood care, carrying of large objects, constructing bridges, foraging, patrolling, chaining, sorting eggs, nest building, and nest grooming. Social insects constitute 2% of insects, half being ants. Besides ants, termites, bees, and wasps have been observed to exhibit some forms of SI behavior as in the aforementioned tasks.

Apart from insects, many other biological groups exhibit behavior with some of the features of SI, such as flocks of birds and schools of fish. A seminal model of artificial flocks and schools of fish was proposed by Craig Reynolds (1987).

It is a computational model for simulating the animation of a group of entities called “boids,” i.e., it is intended to represent the group movement of flocks of birds and fish schools. In this model, each boid makes its own decisions on its movement according to a small number of simple rules that react to the neighboring members in the flock and the environment it can sense. The simple local rules of each boid generate complex global

behaviors of the entire flock. In addition to being used to simulate group motion in a number of movies and games, this flocking behavior has been used, e.g., for time-varying data visualization (Moere 2004). For more examples, see Hassanien and Emary (2016).

In studying these biological systems, several concepts of relevance to SI were recognized. They can be summarized as:

1. Multiple communications (of various types) among units
2. Randomness (random fluctuations)
3. Positive feedback, to reinforce random fluctuations
4. Negative feedback for stabilization

Of the various types of communication, we have already noted *stigmergy*, i.e., *indirect communication* by modification of the environment. On the other hand, *direct communication* may occur *unit to unit*, contact being a special case (e.g., via antennae or mandibles in insects), or by *broadcasting* within a certain range (e.g., acoustically or chemically). The type and specific mode of communication has been found to be critical to the task performed, as, e.g., in what types of patterns are formed (Eftimie et al. 2007).

Finally, the most basic lesson from biological studies of SI is that biology has found solutions to hard computational problems and that the design principles used in doing this can be imitated.

Robotic Systems

The actual realization of SI systems as collections of robots is a very hard problem; in fact, it is quite difficult to make even small groups of robots perform useful tasks (Parker et al. 2005; Sahin and Spears 2005). For a review from an engineering perspective, see, e.g., Brambilla et al. (2013). Making even a single mobile, autonomous robot work in a reliable way (even in simplified environments) is a complex project. Often the technical problems with small groups of robots are quite far from the goal of SI, so there is not much reason to use the term “swarm.” Terms such as

“collective robotics,” “multi-robot systems,” and “distributed autonomous robotic systems” are generally, and more appropriately, used. But, whenever the tasks carried out by these robotic systems become scalable to large numbers, the term “swarm robotics” is appropriate, and, in fact, it has come into use. More typically, “swarm robotics” simply describes the design of groups of robotic units performing a collective task. Each robotic unit cannot solve the task alone; and collectively, the robotic units try to accomplish a common task without centralized control.

As for any robotic system in general, each robotic unit, and the group as a whole, requires design of mechanics, control, and communications. The emphasis of current research, in relation to swarm robotics, is primarily on the latter two: (1) effective *communication* among the robot units (Hamann 2010) and (2) effective *control* via decentralized algorithms and robustness (Sahin et al. 2007). For a recent update on swarm robotics, see, e.g., Hamann (2018). Research in robotic *communication* has become important with the growth of wireless communication networking and the lower cost of building robotic units, thus opening a new range of applications for multi-robot systems with networking capabilities, including swarm robotics. In fact, swarm robotics provides the common ground for convergence of information processing, communication theory, and control theory (Hamann 2010, 2018).

1. Research in *control* of robotic swarms is particularly important to guarantee the stability of the swarm since the swarm does not have a centralized control. The stability of a swarm is a special case of the general problem of distributed control. In fact, after swarm robotics algorithms for task implementation have been devised, the practical realization requires stability and robustness, i.e., proper control. Swarm control presents new challenges to robotics and control engineers: various types of controllers for swarms are currently being investigated, e.g., neural controllers (Sahin et al. 2007). The control theory example

coming closest to the problem of swarm control is perhaps that of “formation” control, e.g., the control of multi-robot teams or autonomous aircrafts or land or water vehicles. These studies, when extended to decentralized systems, lead to consider problems of asynchronous stability of distributed robotic systems and swarms (Gazi and Passino 2011).

Although much progress has been made in swarm robotics, the application of SI algorithms is still underdeveloped; one reason is that often the SI behavior emerges only above a critical number which is too large to make the construction of a robotic swarm practical, because it is too complex or expensive. Investigations of this type are thus generally carried out by simulation (Pinciroli et al. 2012).

These simulations are specialized methods of swarm robotics. Early examples included “executable models” which can run in simulation or on a mobile robotic unit and can execute all aspects of a robotic unit behavior (sensing, information processing, actuation, motion), i.e., they fully represent how perception is translated into action by each robotic unit. Executable models were an evolution of early protocols (the so-called “behavior-based” protocols) designed around the subsumption architecture (Brooks 1986). Behavior-based protocols were generalized into Markov-type methods, i.e., protocols where the transitions between the possible states of a robotic unit are specified by a probability transition matrix as in Markov processes (Johnson et al. 1998). More recently, simulators have improved in flexibility and efficiency via highly modular designs (Pinciroli et al. 2012).

Looking at applications, swarm robotics has by now accumulated a collection of standard problems which recur often in the literature. One group of problems is based on *pattern formation*: aggregation, self-organization into a lattice, deployment of distributed antennas or distributed arrays of sensors, covering of areas, mapping of the environment, deployment of maps, creation of gradients, etc. A second group of problems focuses on *some specific entity in the environment*: finding the source of a chemical plume,

homing, goal searching, foraging, prey retrieval, etc. And a third group of problems deals with more *complex group behavior*: cooperative transport, mining (stick picking), shepherding, flocking, containment of oil spills, etc. This is not an exhaustive list: other generic robotic tasks, such as obstacle avoidance and all terrain navigation, are also swarm robotics tasks.

One envisioned application of swarm robotics which received considerable media attention in the past was the ANTS (autonomic nanotechnology swarm) project by NASA (Curtis et al. 2000). This project envisioned *nanobots* (i.e., a swarm of microscopic robots) operating autonomously to form structures for space exploration.

The European Union-sponsored swarm robotics project (Dorigo et al. 2004; Mondada et al. 2004) was completed in 2005 after demonstrating several critical tasks, such as autonomous self-assembly, cooperative obstacle avoidance, and group transport. For this project, a new type of robot called an s-bot was developed. A swarm-bot could transport an object too heavy for a single s-bot (Mondada et al. 2005). The project has continued to progress and evolved into the “Swarmanoid” project which ended in 2011 (www.Swarmanoid.Org). For another example, see, e.g., Arvin et al. (2014). The largest swarms so far realized are those of the kilobot project. The kilobot swarm is a swarm of 1024 units that can be programmed to experiment with swarm robotics in large-scale autonomous self-organization (Rubenstein et al. 2014).

Although swarm robotics could be defined as the robotic implementation of SI (Definition 1), so far, as noted, this implementation remains a distant goal. Meanwhile, concepts from SI can be usefully applied to collections of cooperating robots. Thus, referring to the intuitive notion of SI (Definition 1), the robotic swarm can be characterized by the type of algorithm and of (decentralized) control, the number of units above which new behavior emerges, the communication method (range, topology, bandwidth), the processing and memory capability of each unit, and the heterogeneity of the group.

Swarm robotics, besides the implementation of SI algorithms, includes the material (mechanical

and electronic) realization of the units comprising the swarm. This is, as noted, an arduous task which often becomes the emphasis of research in swarm robotics. But, as it was emphasized in the early years of SI, even if the material construction of the swarm were accomplished, SI algorithms would remain as the most difficult challenge for swarm robotics. This can be easily seen from the fact that a “robot” swarm with very advanced hardware is already available for experimentation: it is a group of human beings. Each person could be limited in a controlled way, e.g., by allowing each person to handle only a specific device according to specific rules. Algorithms to make such a swarm doing intelligent tasks are in the province of SI, but they are not simple to devise, as common experience shows. Although the notion of “human swarm” has been discussed in the popular press in connection to the development of social media and crowd sourcing, it has not been translated into quantitative algorithms of the SI type (Brabham 2013).

Artificial Life Systems

The areas of *self-organization*, *complexity*, and *artificial life* (or *A-life*) are all older and broader fields than SI and overlap with it to various extents.

A-life is conceptually placed somewhere between science and technology and between biology and robotics. During the mid-1980s, attempts at imitating living systems with machines grew rapidly and resulted in the formation of the research field of “artificial life” (Adami et al. 2012; Aguilar et al. 2014). A-life investigates phenomena characteristic of living systems primarily through computational and (to a lesser extent) robotic methods.

Its scope is wide, ranging from investigations of how lifelike properties develop from inorganic components to how cognitive processes emerge in natural or artificial systems. It includes research on any man-made systems that mimic the characteristics of natural living systems. By this criterion, it includes SI, but actual, current A-life research is not much focused on SI; rather it

focuses on origin and synthesis of life, evolutionary robotics, morphogenesis, learning, etc.

The basic theories at the foundation of A-life, and of relevance to SI, are the theories of *self-organization* and *complexity*. A-life studies systems which are typically characterized by many strongly coupled degrees of freedom. Systems of this type are more generally investigated within the science of *complexity* which began to be an active field of research in the early 1980s. It is multidisciplinary, and it investigates physical, biological, computational, and social science problems, including a vast range of topics (Traub) from environmental sciences to economics as it is clear from the content of this encyclopedia. One basic feature that these systems have in common is the emergence of complex behavior from simple components, a notion we also find in SI.

In regard to self-organization, we note that, as many systems in nature, A-life systems may start disordered and featureless and then spontaneously organize themselves to produce ordered structures, i.e., they self-organize. The theory of self-organization, going back to the 1950s (Nicolis and Prigogine 1977), grew out of a variety of disciplines but mainly from thermodynamics, nonlinear dynamics, and control theory. Self-organization can be defined as the spontaneous creation of a globally coherent (i.e., entropy lowering) pattern out of local interactions – a concept also relevant to SI.

Because of its distributed character, self-organization tends to be robust, resisting perturbations. The dynamics of a self-organizing system is typically nonlinear, because of feedback relations between the components. Positive feedback leads to fast growth, which ends when all components have been absorbed into the new configuration, leaving the system in a stable, negative feedback state. Nonlinear systems have in general several stable states, and this number tends to increase (bifurcate) as an increasing input of energy forces the system away from its thermodynamic equilibrium. To adapt to a changing environment, the system needs a variety of stable states that is large enough to react to perturbations but not so large as to make its evolution

uncontrollably chaotic. The most adequate states are selected according to their fitness, either directly by the environment or by subsystems that have adapted to the environment at an earlier stage.

Formally, the basic mechanism underlying self-organization is the (often driven by randomness) variation which explores different regions in the system's state space until it enters an attractor. This precludes further variation outside the attractor and thus restricts the freedom of the system's components to behave independently. It is equivalent to the decrease of statistical entropy that defines self-organization.

It is useful to keep this brief sketch of self-organization theory in mind as we proceed in describing SI, since the concepts in the theory of SI are evolved from a combination of concepts of self-organization and computation.

Definition of Swarm

After having looked, in the previous three sections, at actual robotic, biological, and A-life systems and ideas of complexity and self-organization related to SI, we can return to the intuitive definition of SI (Definition 1) and make it more quantitative.

The intuitive notion consists of four elements: SI is “intelligence” achieved “collectively,” “randomly,” and “unknowingly.” An elementary swarm retaining these four elements can be defined as:

Definition 2 (Elementary Swarm) An ordered set of N units described by the N components v_i ($i = 1, 2, \dots, N$) of a vector v ; any unit i may update the vector, at any time t_i , using a function f of K_i vector components. $\forall i \in N: v_i(t+1) = f(v_{k \in K(i)}(t))$.

The elementary swarm describes an internally driven “collective” action. External input may be added in the function f . “Randomness” is built in the updating times. The evolution occurs “unknowingly” since the units have no processing capability. The elementary swarm can be generalized so that randomness appears also in the

parameters of the function f . A further generalization is obtained by letting each vector component to be not just one number but a set of parameters.

Hereinafter, we call *Swarm* (capital S) any system capable of SI. It is worth noting that even in Swarms more general than the elementary swarm, the modeling is assumed restricted in such a way that *no unit is capable of computing the Swarm's next global state* (see also section “[Swarms of Intelligent Units](#)”). Finally, “intelligence” is expected to be achieved by running appropriate algorithms via the updating function f . If and how this is going to be possible requires a more mathematical discussion, which is the subject of sections “[Standard-Mathematics Methods](#),” “[Swarm Optimization](#),” “[Nonlinear Differential Equation Methods](#),” “[Limitations of Standard-Mathematics Methods](#),” and “[Cellular-Computing Methods](#).”

Standard-Mathematics Methods

The science of biological swarms and the engineering of robotic swarms, as well as research in A-life relevant to SI, have progressed by using a broad range of mathematical techniques. All these techniques can be classified in two main groups: (1) “*standard-mathematics*” methods and (2) “*cellular-computational*” methods.

By *standard-mathematics methods* (SMm), we mean any method that is based on the standard tools of applied mathematics and computations based on standard (von Neumann) computer architectures. Examples are methods in differential equations, stochastic techniques, linear systems, and optimization. By *cellular-computing methods* (CCm), we mean highly parallel and local computational methods, with simple cells as the basic units of computation, typically carried out on cellular automata (CA) (see “[Cellular Automata](#)”) (Sipper 1999).


These two mathematical approaches reflect two distinct trends in the evolution of SI research, as described below. We consider first (sections “[Swarm Optimization](#),” “[Nonlinear Differential Equation Methods](#),” “[Limitations of Standard-Mathematics Methods](#)”) the approach to SI

based on SMm, since the greatest number of significant results in the area of SI has been obtained, so far, by standard-mathematics methods, specifically in the areas of *optimization* and *nonlinear dynamics*. We consider them in turn in the next two sections.

Swarm Optimization

Optimization is by far the largest research area associated with SI. This is due, mainly, to two extremely successful optimization methods, whose origin is related to models of SI. The two methods are the ant colony optimization (ACO) (Dorigo 1992; Dorigo and Stutzle 2004) and the particle swarm optimization (PSO) (Kennedy et al. 2001; Kennedy and Eberhart 1995). Both ACO and PSO originated in the early 1990s and have resulted in hundreds of applications based on variations of the original algorithms. So much so that the field of “swarm optimization” could stand alone, apart from its relation to SI with which it is sometimes even identified. A thorough and recent description of swarm optimization techniques is in Sun et al. (2011). Here, only the key concepts of swarm optimization are reviewed, as they relate to SI. For a review see, e.g., Nayyar et al. (2018).

In PSO and ACO, as in any optimization methods, a function must be optimized, e.g., minimized. To find the minimum of the function, the variable is changed in a systematic way – the optimization method. Generally, the variable spans a multidimensional space. The search for the global minimum is nontrivial since the function may have many local minima, and the search could end into one of them (see Fig. 2). Various techniques to avoid this trapping have been developed by using some degree of randomness in the search strategy. For example, simulated annealing (Kirkpatrick et al. 1983) was developed to overcome the limitations of nonrandom methods, e.g., the gradient descent (Snyman 2005). PSO and ACO belong to this class of optimization techniques that make use of randomized searches. A recent set of studies on swarm optimization is in Tan et al. (2017).



Swarm Intelligence, Fig. 2 Simplified illustration of the typical problem encountered in optimization. Starting from the value represented by the *open circle* and varying the

parameter continuously the algorithm will find one of the nearest local minima (*black circles*) rather than the global minimum (indicated by the *arrow*)


Particle Swarm Optimization (PSO)

PSO was developed by Kennedy and Eberhart in (1995) inspired by the social behavior of bird flocking and fish schooling (Reynolds July 1987).

In PSO, the position of each unit of the swarm is a point in the variable space of the function to be minimized. Every unit tries to reach the position corresponding to the minimum of the function. Each unit is assumed to know the global minimum value of the function and to detect the value of the function at its location as well as the value of the function at the locations of a group of neighbors. The size of the group of neighbors is a parameter and could be the whole swarm.

The algorithm is, schematically, as follows. The units are initially in random locations. Every unit moves in the variable space and remembers the location L1 where, among the locations visited so far, the function had minimum value. It also remembers the location L2 where, among the locations visited by all its neighbors, the function had minimum value.

At each time step, each unit calculates its distances from L1 and from L2, forms a weighted average of L1 and L2 using random weights, and changes its velocity in proportion to this weighted average (see Fig. 3). As a result, every unit tends to move toward the location of the global minimum by taking advantage of it and its neighbors' knowledge. The process stops either when a unit is sufficiently close (by a chosen tolerance) to the location of the global minimum or when a chosen maximum number of iterations has been run.



Swarm Intelligence, Fig. 3 Illustration of the velocity update mechanism in PSO. A unit (*white circle*) originally moving to the left changes its velocity by adding the two components in the direction of L1 (the unit's best location so far) and of L2 (the neighbors' best location so far). The magnitude of these components is determined by random weights. The new velocity (*heavy solid line*) generally tends to be in the direction of the global optimum

PSO belongs to the category of stochastic, population-based algorithms, such as, e.g., genetic algorithms. The PSO's great merit is its simplicity. In many cases, it outperforms genetic algorithms. Similarly to all optimization algorithms of this type, PSO convergence relies on the use of heuristics; and convergence does not mean convergence to the optimum. In fact, the basic PSO does not guarantee convergence even to a local minimum. The basic PSO is also inefficient in dynamic optimization problems, i.e., problems in which the optimum location changes. However, variations of the basic PSO have been

proven to have improved performance in dynamic problems and to be capable of convergence to local minima.

Many improvements of the PSO basic algorithm have been developed and applied successfully to a wide range of optimization problems: continuous and discrete, constrained and unconstrained, single and multi-objective, and static and dynamic. Specific applications cover just about all areas of applied optimization. The main classes of applications have been in areas such as:

1. Neural networks (training, supervised and unsupervised learning, architecture selection, etc.)
2. Game learning
3. Clustering
4. Design (aircraft wings, antennas, circuits)
5. Scheduling and planning (maintenance, traveling salesman, power transmission, etc.)
6. Controllers (flight path, air temperature, power stabilizers, etc.)
7. Data mining


For more details, see, e.g., Yang et al. (2013). For a recent review, see, e.g., Sengupta et al. (2018). For applications of PSO to statistical regression, see Mohanty (2018).

Ant Colony Optimization (ACO)

The key idea of ACO (Dorigo 1992; Dorigo and Stutzle 2004) is an abstraction and generalization of the Goss et al. (1989) two-path experiment with ants. A first generalization of the two-path problem is finding the shortest path between the starting point S and the final point F when between S and F, there are many possible paths.

The two-path problem can be represented by a graph with three vertices (S, A, F) and three arcs ($S \rightarrow A$, $A \rightarrow F$, $S \rightarrow F$). The short path is $S \rightarrow F$, and the long path is $S \rightarrow A \rightarrow F$.

If we add another vertex B and make the graph complete (i.e., we join with an arc each vertex to every other vertex), we obtain five possible paths: $S \rightarrow F$, $S \rightarrow A \rightarrow F$, $S \rightarrow B \rightarrow F$, $S \rightarrow A \rightarrow B \rightarrow F$,



Swarm Intelligence, Fig. 4 The five possible paths from S to F: $S \rightarrow F$, $S \rightarrow A \rightarrow F$, $S \rightarrow B \rightarrow F$, $S \rightarrow A \rightarrow B \rightarrow F$, and $S \rightarrow B \rightarrow A \rightarrow F$

and $S \rightarrow B \rightarrow A \rightarrow F$ (see Fig. 4). The idea is easily generalized to a complete graph with more vertices. The number of paths increases exponentially, so checking the length of all the possible paths becomes computationally unfeasible for a large enough number of vertices.

A number N of “ants” start at vertex S choosing randomly which vertex to go next. At every new iteration, each ant decides which arc to traverse next, with a probability proportional to the amount of pheromone on the arc relative to the total amount of pheromone on the possible arcs that the ants could choose.

After an ant, k , reaches the destination F, the length of its path L_k is remembered (if an ant reaches the destination via loops, L_k is calculated after removing the loops). The ant retraces exactly the path (without loops) and deposits pheromones on the arcs in proportion to $1/L_k$.

In this way, the marking (with pheromones) of the paths by all the ants modifies the graph so that the probability of any ant taking the shortest path at the next trip from S to F increases. Eventually, all the ants will follow the same path, if the algorithm converges, and the path will be the shortest if the convergence is to the global minimum. As in most optimization methods, this is never guaranteed.

The idea of the original ACO algorithm (Dorigo 1992) adds to the foregoing sketch of the basic model three more elements. First, each arc has an a priori propensity to be traversed (regardless of pheromone content); second, each ant keeps in memory a taboo list of arcs not to be traversed, to avoid loops; and third, the pheromones evaporate at a given rate.

ACO key insight is the application of the concept of stigmergy to stochastic optimization. The ants communicate by modifying the environments (graph) and act probabilistically on the basis of the modified environment.

Many variations of the basic ACO algorithm have been proposed and implemented. The many variations take advantage of specific knowledge about the specific problem, i.e., they use heuristics, e.g., by setting the a priori propensity of traversing an arc or by setting the evaporation rate.

ACO eventually resulted in a metaheuristic which is a strategy for designing ACO heuristics. Various ACO-based metaheuristics have been developed. Similarly to PSO, ACO algorithms have been applied to all the basic types of optimization problems: continuous and discrete, constrained and unconstrained, single and multi-objective, and static and dynamic. The first application of ACO was to the traveling salesman problem, which is an NP-hard combinatorial optimization problem, and it is the most frequently attacked problem using various ACO heuristics.

The main classes of other applications are to problems of:

1. Ordering (scheduling, routing)
2. Assignment (neural network training, image segmentation, design)
3. Subsets finding (maximum independent set)
4. Grouping (clustering, bin packing)

Clearly, “swarm optimization” successfully uses concepts from the general notion of SI, but optimization is not in itself a necessary characteristic of SI. In fact, many tasks actually or potentially carried out by swarms are not optimal in any sense. See also Monmarché (2016).

Nonlinear Differential Equation Methods

One fruitful approach to modeling swarms has been to treat each individual as a discrete particle. These “individual-based” models have been employed in quite a few biological and mathematical studies. They are based on simple rules of motion for each

individual, involving some combination of self-propulsion, random movement, and interaction with neighboring organisms. The models typically take the form of coupled *nonlinear* difference or differential equations, which may be stochastic or deterministic, depending on the particular features of each model. Numerical simulations have revealed collective behavior. But a main disadvantage of such models is that, for realistic numbers of individuals, analytical results for the collective motion are difficult or impossible to obtain. It is worth mentioning that some progress has been made in obtaining analytical results for stationary groups. In Mogilner et al. (2003), a discrete model was formulated, and a Lyapunov functional was used to successfully predict an equilibrium state of equally spaced organisms. However, analytical (nonstatistical) descriptions of non-equilibrium states in discrete swarm models are few.

Other investigations of swarming have been carried out in a continuum setting, in which relevant quantities are described as scalar or vector fields. This approach goes back to 1980; reviews are provided in Murray (2007). Continuum models may be constructed a priori or by coarse graining a particle model. In general, continuum models provide a convenient setting in which to study large populations, since one may apply machinery from the analysis of *partial differential equations*. In the context of swarms, the focus has generally been on models in which the population density satisfies a convection-diffusion equation ensuring that the population density is conserved while individuals travel with a set average velocity. Models of this type (Topaz and Bertozzi 2004) can predict, e.g., whether a population aggregates or disperses, the regions of aggregation, and length scales of the density patterns. Many applications to biological and other collective systems have been carried out recently (Topaz et al. 2012; Canizo et al. 2011), and this branch of swarm studies is also expanding (Elamvazhuthi et al. 2018).

Limitations of Standard-Mathematics Methods

In describing the SI investigations in the previous two sections, we have encountered the concepts of

optimization and *nonlinearity* (and earlier, in section “[Artificial Life Systems](#)” *complexity* and *self-organization* arising from nonlinearity). To see more precisely the relation of these four concepts to SI, we refer back to the definition of elementary swarm (Definition 2). Note that, by this definition, the Swarm is in principle capable of optimization, complexity, and self-organization. In fact, it is clear that a Swarm is a *self-organizing* system, by definition, and that, depending on the choice of the updating function, f , the pattern formed by the swarm components might, in principle, achieve high *complexity*.

As for nonlinearity, it is convenient to think of f as a function of a function.

$$f = f(g(K_i))$$

where g represents the dependence from the neighbors. While g maybe linear or nonlinear, the function f , representing the mode of updating, will typically be nonlinear and/or probabilistic. In spite of this, the evolution of the Swarm only in special cases can be modeled by standard *non-linear dynamics* as studied via nonlinear differential equations (as we have seen in the previous section).

In regard to *optimization*, the elementary swarm can be easily designed to be an optimizer. In fact, if the updating is sequential, and f and v_i are chosen appropriately, one obtains a simple PSO system.

Thus, the elementary swarm describes a simple but powerful system capable, in principle, of self-organizing and producing complex structures and optimal solutions. On the other hand, even the elementary swarm is more general than these properties; a Swarm is not restricted by the notions of optimization or nonlinear dynamics (and self-organization or complexity tied to nonlinear dynamics). All these can be properties of the swarm, but none is a requirement for SI. To find out what SI can do that is beyond what we have described so far, we must look at the computational capabilities of swarms. And for this, we need to look at cellular-computing methods since standard-mathematics methods are ill suited to deal with computation.

Cellular-Computing Methods

Some of the first studies in SI were based on computational models and, more precisely, on distributed algorithms applied to robotic units (Beni 1988, 1992; Beni and Hackwood 1992).

In these computational models, the swarm was developed as an evolution from a distributed system of processors, as follows. In *distributed computing*, the algorithms are designed for a “static” set of processing units, where “static” is meant literally as “not moving.” For illustration, if a set of CPUs, computing in a distributed way, via wireless communication, started moving around, this system would look very much like a robotic swarm. In fact, referring to the intuitive notion of SI (Definition 1), all points of the SI definition would be satisfied by such a dynamic, distributed computing system provided the CPUs had, in some sense, limited capabilities.

The main point is that this distributed computing swarm differs from the robotic, and non-robotic, swarms described in the previous sections (“[Standard-Mathematics Methods](#),” “[Swarm Optimization](#),” “[Nonlinear Differential Equation Methods](#),” “[Limitations of Standard-Mathematics Methods](#)”) in that the *intelligent task of the swarm is now seen as a “computation.”* And this focus on computation leads us now to consider the other broad set of techniques used in SI research, i.e., techniques based not on standard mathematics but on *cellular computing*.

Cellular computing differs qualitatively from the standard von Neumann computing architecture. The latter is based on one complex processor that sequentially performs, at each time step, a single complex task. In contrast, in *cellular computing*, a very large number of simple processors (cells) are the units of computation. They compute (typically) in parallel with local connections between cells. The qualification “simple” can be made precise by requiring, e.g., each cell to be a “finite state” machine. Cellular automata are the most obvious examples of cellular-computing systems, but cellular computing applies to many other systems as well (Sipper 1997).

By definition, then, cellular computing contains several of the features of SI. The elementary

swarm of Definition 2 can be regarded as performing a cellular computation. And in fact, cellular computing has been used extensively in A-life studies, including systems with strong relation to SI (Adami et al. 2012). Cellular-computing systems offer SI something that the SI systems described in the “standard-mathematics” sections (“[Standard-Mathematics Methods](#),” “[Swarm Optimization](#),” “[Nonlinear Differential Equation Methods](#),” “[Limitations of Standard-Mathematics Methods](#)”) lack, i.e., a clear characterization of *intelligent task*.

Intelligence as Universal Computation

Intelligence is an ambiguous concept, escaping a unique definition (Gottfredson 1997; Legg and Hutter 2007). By identifying “intelligence” with “computation,” the concept is restricted, but, at the same time, it can be made precise. In fact, in SI, we define intelligence unambiguously as the *ability to carry out universal computation*.

Universal computation (or universality) is the property of a computer system (or language) which, with appropriate programming, can be made to perform exactly the same set of tasks as any other computer system (or language).

Universal computation (i.e., the ability to emulate a universal computer) is essentially the limit of any model of computation (Church-Turing thesis) (Cooper 2003). It was first proven by Turing in 1936 that no system can ever carry out explicit computations more sophisticated than those carried out by a Turing machine. Subsequently, universality has been found to be a widespread property of many cellular-computing systems (Wolfram 2002).

One of the first cellular-computing systems shown to be capable of universal computation is Conway’s game of life (Gardner 1970). This CA is also the prototypical example of A-life systems. And it is also an example of the strong connection between universal cellular-computing and bioinspired systems.

More recently, a large number of simple cellular-computing systems have been found to be capable of universal computation (Langton

1984; Dennunzio et al. 2012). Many of these systems are CA, or related systems, using very simple rules of evolution with local interactions. And so they are useful starting points for modeling SI.

In particular, cellular computing is the most appropriate to endow the swarm with the property of *unpredictability*. The latter property was an original motivation for SI (Beni and Wang 1989a, b), and it is crucial in the task of escaping detection by a predator; it is also of importance in engineering swarms for strategic defense applications.

Unpredictability is almost a built-in property of cellular-computing systems because if one observes the rules of evolution in their raw form, it is usually almost impossible to tell much about the overall behavior they will produce. See, e.g., Agapie et al. (2014).

Relations to Standard-Mathematics Methods

SMm cannot provide the swarm with the element of universal computation, which we have taken as the working definition of “intelligence.” The only way would be to make each unit a von Neumann (i.e., standard) computing system. In a sense, this violates the notion of Swarm, since in a Swarm, by definition, each unit must be “simple.” (This point will be further clarified in section “[Swarms of Intelligent Units](#)”.) On the contrary, the main advantage of cellular-computing systems over standard-mathematics systems is the possibility of universal computation by simple units.

For this reason, CCm are the natural paradigm for the understanding and designing of SI systems, in spite of the fact that the approach to SI based on cellular computing has so far produced fewer so-called SI applications than the approach based on SMm. Indeed, SMm have basic limitations for modeling SI. This is because the use of SMm tends to restrict the range of tasks performable by the Swarm. And this happens because SMm typically solves problems by specifying *constraints*, i.e., conditions to be satisfied by the solution, e.g., by specifying equations. But most computational problems cannot be solved in this way.

The optimization methods described in the previous sections (PSO, ACO, etc.) illustrate the point. In these iterative methods, the key issue is what kind of changes should be made at each iteration step. Starting from a random pattern, at each step, a change is made to get the pattern closer to satisfying the constraint(s). Since direct methods (e.g., gradient descent) rarely work as the pattern gets stuck into local minima, randomness in updating is added. In this way, larger portions of the solution space are sampled. The larger the changes made, the faster one can potentially approach a global minimum but the greater the chance of overshooting. The result is that no iteration technique of this type can guarantee a solution to general combinatorial optimization problems. As we have seen, the swarm optimization methods (e.g., ACO, PSO) rely on heuristics to adjust the search and obtain (nonoptimal but) often satisfactory solutions. But, in general, for the great majority of combinatorial optimization problems (e.g., the traveling salesman problem Johnson and McGeoch 1997), no polynomial upper bound on the time complexity has been found so far. And this happens in many problems whose solution is sought by using randomness to satisfy the imposed constraints. As an example, a set of identical balls cannot be shaken into an ordered, closed-packed configuration. With extremely high probability, they lock into some configuration or another, not the optimal (close-packing) one.

This fact has *important implications* for SI. What it says is that no matter how much randomness is added to the system, it may never evolve to reach the solution specified by the constraints. Although, ultimately, constraints can be set up as a way of specifying algorithms, and hence computing, it is far simpler to specify algorithms via rules of evolution, as it is done in cellular computing.

The conclusion is that methods based on constraints and other SMm are not ideally suited for systems evolving with great complexity, and, in particular, they are not suitable for universal computation. Thus, if SI is to be a framework for (biological or engineered) swarms to carry out “intelligent tasks” with the greatest generality, a methodology that allows for the swarm to carry

out universal computation is necessary. To this aim, CCm are the most suitable.

Unfortunately, although CCm have many advantages over SMm for modeling SI, they address only three of the four key elements of the notion (Definition 1) of SI (“intelligence,” “collectively,” “randomly,” “unknowingly”), leaving out the element of “randomness.” Generally, CCm operate deterministically and do not include “randomness,” as, e.g., “swarm optimization” systems do. But this does not have necessarily to be the case. The issue is addressed in the next section.

Randomness in Swarm Intelligence

Randomness is a key element in the notion of SI (cf. Definitions 1 and 2). Examples from biology justify this requirement. Randomness is not easily quantified precisely, but, whatever the form and measure chosen, the point is that for swarms, some form of randomness is necessary – otherwise, they would fail to be models for analyzing a large class of biological systems. But what kind of randomness is essential to model these biological systems?

Randomness in the number and type of agents is not important – the agents could be strictly identical and remain in the same number. Randomness in the initial conditions is not essential either. Many swarms evolve from regular initial conditions into highly complex and random patterns. Randomness of external input from the environment is not always present, and it is certainly not a requirement for biological swarm behavior.

What about the randomness artificially added to the units, as in swarm optimization?

The randomness added to the units in PSO or ACO algorithms is modeled as originating from the random behavior of each unit. This is a plausible assumption in relation to biological systems. But the swarms in PSO and ACO are typically updated in an orderly (nonrandom) way, typically sequentially (there are also parallel implementations Olariu and Zomaya 2005), whereas, in biological systems, the units update in a disordered, random fashion.

And it is this type of randomness that is both *necessary* in any biologically relevant model of

swarms and *sufficient* to provide many (but not all) of the advantages of randomness in solving swarm engineering problems.

The conclusion is that *the only randomness that is truly essential for SI is randomness in the times of operation of the units*. Each unit has its own clock, not synchronized with other units' clocks. Other types of randomness in the behavior of the units or the environment may be required to solve specific problems, but randomness in times of operation is necessary for any biologically realistic model.

Interestingly though, many applications so far considered in the area of SI do not yet include this randomness in the models. We have already mentioned that typical optimizing swarms update sequentially and CA systems operate largely in parallel, i.e., synchronously. Synchronous or sequential operations are by far the most common updating modes in either SMm or CCm.

The Implicit Assumption of Asynchrony Irrelevance

As noted, it is a basic fact that biological agents, apart from exceptional cases, do not operate synchronously (or sequentially) in groups. It is also a fact that people in social groups do not operate synchronously or sequentially. If SI is supposed to model biological and social swarms, SI must be based on models that do not operate synchronously or sequentially (Huberman and Glance 1993).

And if biological swarms are capable of solving problems (including optimization) without synchrony (nor sequentially), as they do, then models that imitate those swarms should operate asynchronously (not sequentially).

But, as noted, the main modeling paradigms for bioinspired algorithms, standard mathematics, and cellular computing are either essentially sequential or synchronous.

An example from SMm is the solution of partial differential equations: they operate synchronously on every point (clearly seen in solving them numerically and iteratively). This unrealistic use of differential equations in biological

processes has been pointed out, e.g., in the problem of morphogenesis (Liang and Beni 1995). The Turing diffusion-reaction model (Turing 1952), being based on differential equations, implies synchronicity and central control; hence, it is physically not realistic for a scale of the order of 100 cells.

In fact, synchronicity leads to realistic models only whenever the spatiotemporal resolution is high, as, e.g., for phenomena typically studied in physics. But when the units studied are complex or few enough to have a less fine spatiotemporal resolution, as in biology or human societies, synchronicity is not realistic, as it is obvious by observation.

Thus, in using synchronous methods for biological or human societies, implicitly a *strong assumption* is being made, i.e., *that the synchronously (or sequentially) and nonsynchronously (and not sequentially) obtained solutions would coincide*.

But *this assumption has no validity*. In fact, it has been shown, for example, that CA, when running in synchronous and nonsynchronous ways, normally produce totally different results (Cornforth et al. 2005). This has been noted already in the 1990s (Bersini and Detour 1994) in A-life studies. In Bersini and Detour (1994), two well-known CA were compared: Conway's "game of life" (Gardner 1970) and the immune network model. The former is a two-dimensional CA capable of universal computation when run synchronously. But the behavior is totally different when run without synchrony: the game of life stops producing complex patterns and converges to a fixed point.

The immune network model is asynchronous and the game of life synchronous. The crucial factor in the different behavior of the two systems was identified as the synchronous versus asynchronous updating. In fact, it was concluded that, in this case, asynchrony induces stability in CA. This agrees qualitatively with studies in standard mathematics (Bersini and Detour 1994).

In conclusion, the assumption that asynchrony makes no difference has been found not to be valid. (for an example in PSO see, e.g., Nor Azlina et al. 2014) Hence, asynchronous systems must be

studied as such, not by using synchronous models. Moreover, different types of asynchrony yield different results, as discussed in section “[Asynchronous Swarms](#).”

Asynchronous Swarms

Several cellular-computing studies in the 1990s (Sipper 1999; Schonfisch and de Roos 1999) led to a variety of results emphasizing the role that different types of asynchrony play in the results. Studying asynchronous systems is complicated because, among other things, deviation from synchronicity, i.e., from the mode of updating all units in parallel at each time step, may occur in several different ways. For example, sequential updating and random updating are both asynchronous but very different.

Types of Asynchrony

Unfortunately, there is no standard vocabulary for the various types of asynchrony. Thus, we use the following classification to describe the possible types of asynchrony.

Consider an updating cycle (UC), i.e., the time interval at the end of which all units have been updated at least once. Eight types of UC can be identified by the presence or absence of any of the following three properties: *synchronicity* (S), more than one unit may update at each time step; *multiplicity* (M), any unit may update more than once in each UC; and *randomness* (R), the updating order varies randomly for every UC (see Fig. 5).

Four units are represented by rectangles with different patterns, from black (bottom) to white (top). The horizontal axis measures time steps in units equal to the base of a rectangle. The vertical dashed line indicates the end of an updating cycle (i.e., all units have updated at least once). The label below the horizontal axis specifies the type of updating ($\sim$ means “not”). The standard “parallel” and “sequential” updating are, respectively, (S, $\sim$ M, $\sim$ R), Fig. 5e, and ($\sim$ S, $\sim$ M, $\sim$ R), Fig. 5a.

These eight basic types of asynchronous updating can be further specialized. For example, if all three properties are absent ($\sim$ S, $\sim$ M, $\sim$ R), the updating is sequential. But, the sequential updating order of the units can be fixed in different ways. Studies of CA have proven that the behavior differs markedly not only for the eight types of asynchrony but even among different sequential ordering (Sipper 1999; Cornforth et al. 2005).


In Wolfram (2002), the (S, M, $\sim$ R) form of updating has been applied to describe processes where each unit has independent clocks; but the clocks have a fixed, nonrandom frequency. This type of asynchrony is considered a good model for forest ecosystems, fire spread, and other natural and artificial systems. The results are very different when updating of the type ($\sim$ S, M, R), ($\sim$ S, $\sim$ M, R), or sequential ($\sim$ S, $\sim$ M, $\sim$ R) is applied to the same system.

In conclusion, the crucial point is that (Cornforth et al. 2005) *the exact manner of updating can have a profound effect on overall system behavior*. The implication of this is that when comparing models of natural systems or artificial multi-agent systems, it must be stated which updating scheme has been used; otherwise, meaningful comparison between different studies may not be possible.

In particular, returning to swarm optimization, one may ask whether in tasks, such as finding the shortest path, it is realistic to apply to natural systems (such as insect societies) swarm models which are sequential (such as swarm optimization models) or synchronous (such as models based on differential equations). While these models work effectively as artificial swarms, there is no proof that they apply to natural systems, which are asynchronous.

Modeling Asynchrony by Synchronous Swarms

Because of the widespread use of synchronous methods in simulations of SI, one might wonder under what conditions a synchronous but stochastic model could be equivalent to an asynchronous one.



Swarm Intelligence, Fig. 5 Illustration of the eight types of updating, according to synchronicity, multiplicity, and randomness. (a) Asynchrony of type ($\sim S, \sim M, \sim R$). (b) Asynchrony of type ($\sim S, \sim M, R$). (c) Asynchrony of type

($\sim S, M, \sim R$). (d) Asynchrony of type ($\sim S, M, R$). (e) Asynchrony of type ($S, \sim M, \sim R$). (f) Asynchrony of type ($S, \sim M, R$). (g) Asynchrony of type ($S, M, \sim R$). (h) Asynchrony of type (S, M, R)

To answer this question, let us consider the two types of stochastic models most commonly used to model randomness in synchronously updated systems. The randomness may be included in (1) the possible outcomes of the updating function or (2) in the choice of the function applied to the updating.

Referring to the definition of elementary swarm (Definition 2), the two cases correspond to generalizing the updating function as follows:

- Case (1) $\forall i \in N: v_i(t+1) = f(v_{k \in K(i)}(t), \zeta)$, where ζ is a random variable.

- Case (2) $\forall i \in N: v_i(t+1) = f_{(t)}(v_{k \in K(i)}(t))$, where $P[f_{(t)} = f_\gamma]$ is the probability mass function of choosing $f_{(t)} = f_\gamma$ out of a set of N_f possible functions $\{f_\gamma; \gamma = 1, \dots, N_f\}$.

Case (1) is typical of probabilistic CA, and it is also the method used in PSO. In these systems, the state vector, at each time step, evolves according to a fixed rule which produces a new state vector from the previous one. The rule is based on the state of the neighbors of each unit and does not change from step to step, but the outcome of the rule is probabilistic.

Case (2) is what is done, for example, in probabilistic iterated function systems (Peruggia 1993). In probabilistic iterated function systems, a vector evolves via a set of maps (a map is a function whose domain and range coincide); at each time step, a map is chosen, probabilistically, from a set of possible maps.

In either cases (1) or (2), the updating scheme fails to model the actual time evolution of natural systems not so much because the updating is applied synchronously but because the randomness is applied *collectively*, i.e., to all the units in the same way. On the other hand, a synchronous algorithm realistically simulating independent random updating can be run as case (2) applied *individually* to each unit, as follows:

- Case (3) $\forall i \in N: v_i(t+1) = f_{(t)i}(v_{k \in K(i)}(t))$, where $P[f_{(t)1} = f_{\gamma 1}, f_{(t)2} = f_{\gamma 2}, \dots, f_{(t)N} = f_{\gamma N}]$ with $f_{(t)i} \in \{f_{\gamma}; \gamma = 1, \dots, N_f\}$ is the joint probability mass function of each unit i updating, at time t , according to the function $f_{(t)i}$.

In the simplest embodiment of case (3), the set of possible updating functions consists only of the identity and of another function f , with probabilities p and $(1-p)$, respectively (i.e., a Bernoulli process). In such a case, every unit, at each time step, either does not update, with probability p , or updates according to the function f , with probability $(1-p)$. Running this algorithm synchronously is equivalent to asynchronous independent updating of the units in a random way – a realistic description of a random swarm. So, under these independently stochastic conditions, running a simulation synchronously represents correctly the physical asynchronous updating of the swarm units. On the other hand, this does not change the fact that different results are obtained when using this random updating (whether simulated with stochastic synchrony or not) instead of synchronous or sequential updating.

Local Synchrony and Self-Synchronization

Another approach to dealing with the problem of random updating by the swarm units is to explore

the possibility of self-synchronization. If the swarm can self-synchronize, then all the results for synchronous swarms could be applied.

To look into this, let us return to the classification of the types of asynchrony, i.e., the SMR classification above. If the SMR properties are applied to blocks of units rather than individual units, the resulting updating orders are referred to as *locally synchronous*.

CA with cells organized into blocks have been investigated (Sipper et al. 1997). These CA relax the normal requirement of all cells having the same update rule. Cells within a block are updated synchronously, but blocks are updated asynchronously. They experimented with different SMR types of asynchrony and concluded that synchronous and asynchronous CA can be evolved with equivalent computational properties, but CA of the asynchronous type may require a larger number of cells (Sipper et al. 1997). Another study (Clapham 2002) has shown cases in which local synchronization can lead to the same outcome as with global synchronization. But how can local synchronization be achieved?

A number of schemes have appeared in which the order of updating depends on local interactions and leads to local synchronization. In effect, what local interactions (or constraints) can do is to force a unit to wait to update until others are ready, and so this creates a local synchronization. An asynchronous CA model that can behave as a synchronous CA has been demonstrated (Nehaniv 2002); it functions by the addition of extra constraints on the order of updating, effectively providing a type of local synchronization. Whether these methods of self-synchronization may in some cases result in realistic models of natural systems of SI remains an open question.

The Natural Asynchrony of Swarms

We have seen that the implicit assumption of asynchrony-synchrony equivalence must be rejected and that different types of asynchrony give different results. But what type of asynchrony is most relevant to SI? There is no easy answer. For example, ants work and rest; active

and resting periods have an aperiodic pattern for individual ants, but for the whole colony, there are synchronized periodic patterns of active and resting periods.

In spite of the difficulty of finding a clear-cut answer to the question of the natural mode of SI updating, from observations of biological systems and from local synchronization models, it may be plausible to assume that the essential form of asynchrony in SI is the randomness in the working of the individual clocks, as argued in section “[Randomness in Swarm Intelligence](#)”; hence, the SI asynchrony must be characterized by the presence of all three asynchrony properties, i.e., SMR.

In conclusion, at this stage of our discourse, the Swarm remains defined as in Definition 2, qualified by SMR asynchronous updating, which hereinafter we call *natural* asynchrony. Note that stochastic synchronous simulations of this model can also be carried out as, e.g., in Case (3) above.

The Realization of Asynchronous Swarms

So far, we have established the importance and type of asynchronous models in SI, but what SI investigations using asynchronous swarms are there?

As noted in section “[Asynchronous Swarms](#),” research in asynchronous models is still very limited, relative to synchronous models, and this in spite of the fact that the very first models of SI were all asynchronous, using SMm based on finite differences (Beni 1992). Explicit updating schemes in finite difference methods can also be regarded as parallel CA, thus belonging to both SMm and CCm. Investigations of asynchronicity in finite difference methods are not common (Beni 2004b). Examples include a nonlinear updating rule was based on a linear relation between two neighboring units (Beni and Hackwood 1992). A gradient type of swarm updating was also proposed in modeling morphogenesis (Liang and Beni 1995).

For swarms updating with “natural” asynchrony, i.e., according to SMR, a study (Beni 2004b) gives a proof of convergence to the same

state as by using synchronous or sequential iterations. It was also shown that, under certain conditions, the SMR asynchronous updating leads to convergence while synchronous updating does not. This is another example of the advantages of randomness in allowing the swarm to reach a fixed state.

At the end of section “[Cellular-Computing Methods](#),” we concluded that CCm have, for SI modeling, many advantages over SMm. The most crucial advantage is the possibility of universal computation which we took as the definition of intelligence for SI. We also noted, however, that studies based on CCm which include randomness are scarce. We described a few in section “[Randomness in Swarm Intelligence](#)” especially in discussing the qualitative differences with synchronous CA and in relation to mechanisms of local and self-synchronization.

Generally, these studies model relatively trivial phenomena but cannot model nontrivial phenomena such as universal computation. In fact, there are very few studies of universal computation in asynchronous CA. Significant advances have been made only recently. The first attempts were made by simulating a synchronous CA on an asynchronous CA (Nakamura 1974) after which a synchronous model, as a Turing machine, was simulated on the synchronous CA. However, this asynchronous CA is, in practical realization, synchronous.

Improved asynchronous CA do not rely on global synchronization but conduct asynchronous computation directly by simulating delay-insensitive circuits, i.e., circuits in which delays of signals do not affect the correctness of the circuit operation (Lee et al. 2004). This method essentially uses local synchronization with undetermined exact timing between transitions. In this way, an asynchronous CA, with a hexagonal cell structure, capable of universal computing has been realized (Adachi et al. 2004). Although relying on local synchronization, this type of asynchronous CA can mimic natural phenomena as, e.g., phenomena that rely on chemical reactions which occur only when the right molecules are available in the right positions at the right times.

A computation-universal and construction-universal asynchronous CA have been designed (Langton 1984) and used to implement self-reproducing (von Neumann 1966; Langton 1984) machines. Besides computational universality, construction universality is important in SI because it allows the swarm to be hardware reconfigurable, an important characteristic of many biological systems.

We note that the recent interest in asynchronous CA stems not directly from SI but from nanotechnology. In fact, nanocomputer architectures with asynchronous updating may reduce heat dissipation, an important limiting factor in scaling down the size of computing chips (Lee et al. 2004). In this respect, it is likely that SI concepts will play a major role in nanoscale systems.

More recently (Dennunzio et al. 2012) fully asynchronous CA have been studied for simulation of Turing machines and thus for universality. A family of asynchronous CA has been shown to simulate any Turing machine with any input if the cells are updated infinitely many times with certain updating sequences, in particular, random walk sequences. The computational costs have also been investigated and in some cases found to be reasonable (quadratic). Even more recently a Brownian cellular automaton capable of universal computation has been designed (Xu et al. 2019).

In conclusion, the recent realization (Langton 1984; Dennunzio et al. 2012; Xu et al. 2019) of universal asynchronous CA is a promising major step toward the possible realization of true SI. See, e.g., Fatès (2018). For modeling of biological pattern formation using cellular automata, see Deutsch and Dormann (2018).

Characteristics of Swarm Intelligence

The demonstration of universal computation in asynchronous CA amounts to a validation of the concept of SI. In fact, we can now combine universal computation with the elementary swarm definition (Definition 2) to quantify the intuitive definition of SI (Definition 1), as:

Definition 3 SI is the study of *universal cellular-computing systems updating with natural asynchrony*.

Here, “natural” means SMR asynchrony (see section “Asynchronous Swarms”) or updating randomly in parallel, as in Case (3) of section “Asynchronous Swarms.” We call it “natural” since it models the natural mode of updating of typical biological swarms and human societies.

A few remarks about Definition 3, the four elements of the intuitive notions (Definition 1) are made precise by Definition 3: “collectively” and “unknowingly” are inherent in the structure of cellular computing, “intelligence” is in universal computation, and “randomness” is in the natural asynchronous operation.

Definition 3 deals with CCm and may appear to exclude SMm in SI; but this is not the case. In fact, many SMm, as used in SI, can be regarded as special cases of cellular-computing methods, as, e.g., are swarm optimization and iterative methods in finite differences.

Also, Definition 3 does imply that every SI system must be capable of universal computation; what Definition 3 does is to establish a focus of attention for the SI area of studies and at the same time give a precise and realistic meaning to the kind of “intelligence” aimed at in SI, rather than the often vague and exaggerated meanings given in the popular literature.

Definition 3 also indicates how SI becomes of relevance beyond biological systems and robotics. In fact, SI will likely be an important concept in the future of computation. At very small scales, time delays between computational components cannot guarantee synchrony; the various components must have independent clocks, thus beginning to resemble the operation of a biological swarm. So, swarms are likely to be studied extensively in connections with nanocomputing.

The above arguments bring up the question as to whether SI is nothing more than asynchronous cellular computing. The answer is that designs of asynchronous CA, as investigated in computer engineering, are generally not models of SI. A first, basic reason is that, with possibly few recent exceptions (Takada et al. 2006; Dennunzio et al.

2012; Xu et al. 2019), asynchronous CA do not update “naturally.” A second, more fundamental reason is that in most SI studies (both in natural and technological systems), *the units are dynamic*. When the units of an asynchronous CA are made mobile, a different and more complex set of problems need to be solved due to the changing neighborhoods of each cell. These issues of dynamic reconfigurations of cells have not been addressed in asynchronous cellular-computing designs and are likely to remain outside the scope of research aimed at improved computer architectures. The computational problems arising from dynamically reconfiguring cells are central in SI. We address this issue next.

Dynamics in Swarm Intelligence

In the definitions of SI and Swarm given so far (Definitions 1, 2, and 3), there has been no mention of the dynamics of the units. But a general characteristic of the units of a swarm is that almost invariably they are mobile. In fact, we have already discussed the dynamic nature of swarms in sections “[Biological Systems](#)” and “[Robotic Systems](#)” in relation to biological and robotic swarms. The reason we have so far omitted this dynamic character of the units from the progressively more precise definitions of SI is simply for clarity of exposition: if the dynamics is introduced after all the other elements of SI have been defined and quantified, it is easier to single out its real importance.

From Definition 3, with appropriate specializations, all aspects of SI considered so far can be included in a common core of studies. The most general notion of SI is in fact that of *universal computation carried out with “natural asynchrony” by a cellular-computing system*. But we should add, *whose cells are, in general, mobile units*.

The latter qualification would be unnecessary if the description of the dynamic state could be included among the state variables of the cell. But this is not always possible in computing cells since the computation depends on neighbors that change their locations. The fact that *dynamic*

cellular-computing systems (also called *cellular robotic systems*) are not equivalent to (and very hard to simulate by) cellular computers has been emphasized since the very beginning of SI studies (Beni 1988). Thus, we conclude by stating a definition of SI which, while remaining grounded in the intuitive ideas (Definition 1), includes all the concepts discussed quantitatively in this entry.

Definition 4 SI is the capability of universal computation carried out with “natural” asynchrony by a dynamic cellular-computing system.

As we have seen, studies in the area of SI so far have been concerned with models of collective behavior which, to some limited degree, approach SI as defined above. Even though, to date, no system with SI (Definition 4) has been built (or designed), significant progress has been made, and, from what we have seen, it is reasonable to expect that it can be done.

In fact, although there is not yet proof of universal computation carried out with “natural” asynchrony by a *cellular robotic system* (i.e., a dynamic cellular-computing system), the recent proofs (Takada et al. 2006; Dennunzio et al. 2012; Xu et al. 2019) for “static” cellular computers indicate that this may be possible in the near future. Other future perspectives for SI are discussed in section “[Future Directions](#).”

Unpredictability in Swarm Intelligence

We are now in a position to consider an aspect of SI, which has been inherent to the concept of SI from its inception, i.e., the unpredictability (Beni and Wang 1989a, b) of the Swarm.

The unpredictability of the Swarm agrees with the common intuition that it is usually difficult to predict what a program will do by reading its code and the more so the lower the level of the language used by the program. More precisely, a Swarm, like other universal computers, may be impossible to predict in the sense that even if one knows the rules of evolution and an initial state, it can still take an irreducible (Wolfram 2002, 1985) amount

of computation to actually predict future states. Furthermore, the unpredictability of the Swarm is of a more general character than that of any universal computer because of the randomness inherent in its evolution and because of its dynamics.

Generally a system is (at least partially) predictable if at present time, one of its future states can be known via a computation shorter than the one performed by the system to reach that future state. This is not possible for many types of asynchronous CA. This notion of predictability considers computational time complexity but not actual time. The latter is affected by the actual construction of the computing system. As such it is sensitive also to the dynamic of the system.

In fact, the unpredictability of a Swarm by a von Neumann universal computer has been argued in Beni and Wang (1989a) and Beni (2004a) on the basis of its dynamics. The unpredictability of a Swarm by a cellular automaton has also been discussed (Beni and Wang 1989a; Beni 2004a). Although unpredictability is difficult to quantify when including actual system construction parameters, it is often engineered by adding randomness to the system as in camouflage and cryptography. Also, in animals, randomness and dynamics are methods used by a herd to avoid predators by becoming unpredictable. And so in team sports, such as soccer, unpredictability by the opponent is usually achieved by a combination of randomness and dynamics.

Therefore, although far from proving it, we may intuitively conjecture that among systems capable of universal computation, a Swarm, because it computes universally, but asynchronously (with randomness) and dynamically, could be designed to be *the least predictable*.

Swarms of Intelligent Units

Let us now consider collections of intelligent units, i.e., such that each unit is capable of universal computation. These seem excluded from the definitions of SI given so far – a key characteristic of SI is that intelligence is an emergent property, happening only above a certain critical number of units, and not a property of any of the individual

units. On the other hand, some of these groups are often included in broad considerations of SI (Kennedy et al. 2001) as applied to human societies. Under what conditions can these groups be regarded as swarms?

A simple answer to this question runs as follows. Consider the special case of Definition 4, when the cells are not finite-state machines but universal computing units. As long as the task at hand cannot be accomplished by a single unit, but only by more than a critical number of units, the system operates as a swarm. That this can be the case is supported intuitively and from computational considerations, as follows.

Intuitively, we may refer back to the example of the free market model mentioned as illustration of SI at the beginning of the entry. Each individual contributes only as a trader; the “computation” of the market price is done by the swarm collectively and could not be done by any individual agent. The point is that each unit, albeit “intelligent,” uses only a fraction of its capability, i.e., the trading ability, thus operating, effectively, in a restricted, “non-intelligent” capacity.

Computationally, we have noted in discussing unpredictability that, in spite of theoretical computational equivalence among universal computers, the capability for one universal computer to predict another is limited. And, in fact, at the end of section “[Swarms of Intelligent Units](#),” we put forth the conjecture that a Swarm could be the least predictable universal computer.

But even if this conjecture was not true, it still makes sense to think of a case when a Swarm of universal computers is unpredictable by any of the units comprising it, as it has been argued for the case of units capable of universal computation with von Neumann architecture (Beni and Wang 1989a; Beni 2004a). In this sense, we may think of a swarm of universal computers as being more capable than any one of its units, in spite of the fact that the Swarm and any of its units are computationally equivalent. For these reasons, it makes sense to apply, under appropriate conditions, the notion of SI also to human societies.

We may give then a more complete definition of Swarm by adding the characteristic of unpredictability of the Swarm by its units.

Definition 5 SI is the capability of universal computation carried out with “natural” asynchrony by a dynamic cellular-computing system, none of whose cells can predict the computation done by the Swarm.

The latter specification is obviously redundant for common cellular-computing systems, but it is useful to exclude from SI trivial cases of human social activities, for example, activities of human groups whose association cannot be proven to solve a problem that could not have been solved by any individual alone.

Many popular interpretations of SI have appropriated the label SI to refer to almost any trivial human group activity, such as brainstorming. In such activities, often, there is no way of proving that the final output of the group could not have been predicted by one of the members of the group. Ultimately, however, these considerations require an understanding of the relation between human thinking and computation and thus fall beyond the scope of this entry.

Similarly, beyond the scope of this entry and SI fall studies of multi-agent systems in artificial intelligence (Weiss 2000). In the artificial intelligence area, the emphasis on multi-agent systems is in finding decision algorithms, i.e., “agents,” for open environments in which these agents must operate robustly and rapidly, i.e., “intelligently.” Generally, however, the problem of collective decision-making, organization theory, distributed reasoning, and distributed artificial intelligence is typically beyond the scope of SI. See, e.g., Agrawal et al. (2018).

SI deals with human groups only when they operate at a low level of intelligence. Definition 5 includes human groups of individuals that operate under restrictions which the Swarm can overcome. An example is that of groups of individuals, each with a limited computation device connected wirelessly to neighbors as in a cellular-computing system. With appropriate algorithms, such a Swarm could compute universally, while the individual device cannot. This concept could be applicable in, e.g., defense operations or emergency mass evacuations strategies.

In conclusion, Definition 5 embodies the intuitive definition of SI (Definition 1) and indicates why SI methods can be used to solve problems not solvable by traditional methods. Besides having all the properties of universal CA, the Swarm operates, as natural systems do, by independent clocks with no centralization and can be designed to be dynamic and unpredictable by any system including any of its own units.

Future Directions

The field of SI is only 30 years old, and it is in a formative phase, with SI researchers engaged in a broad range of disciplines. During these formative years, the “conversation” about SI has followed several strands around some common themes with various emphases ranging from speculative inquiries to practical interests, as we have seen in this entry.

An example of a recent direction is the proposed perspective of Swarm cognition (Trianni et al. 2011). This approach combines neurosciences and ethology with swarm intelligence methods. The basic assumption is that cognition is at the core a collective phenomenon generated by interacting elementary units. Such units could be elements of a single individual, as in a brain, but also agents in of a society of individuals, as in an ant colony. Thus, the meaning of cognition is broadened beyond the standard neurological concept in the hope of gaining insights, not only on swarm intelligence but also on standard cognitive processes.

Indeed the meaning of the term SI has tended to broaden, covering now many areas, apparently only weakly related by the same intuitive notions. We have seen, however, how all these SI ideas have a “center of attraction” in a basic concept of SI that can be made precise and quantified (Definition 4 or 5) and thus used to provide unity, continuity, and boundaries, thus preventing the area from broadening to the point of being unable to sustain an effective research community.

With this perspective, SI can also be seen as having an ultimate theoretical goal for the

practical realization of engineered Swarms, whether robotic, biological, or simply computational.

Practically, the goal of SI will remain twofold: to provide models to explain biological societies and to engineer algorithms and devices with capability beyond those of traditional technologies. It will continue to include Swarm robotics and bioinspired algorithms such as swarm optimization methods.

But although commonly regarded as a typical example of bioinspired technology, SI applications are likely to go beyond bioinspired systems. We can see this if we consider how nature-inspired technologies have evolved.

Science discovers laws of nature, and technology makes inventions using those laws, often together with design ideas also derived from nature. Thus, for example, laws of physics and designs inspired by crystal structures are now applied to nanotechnologies; similarly, laws of biology and designs inspired by genetic configurations are applied to make artificial organisms in biotechnology. New attempts are also being made (see, e.g., Wolfram 2002) at discovering laws of computing machines as though they were natural systems, and these discoveries are likely to be used to invent new software algorithms and hardware implementations of those algorithms. SI, besides being bioinspired, can be said to be inspired also by this new science, which, in some respects, can be more general than biology.

In this perspective, SI will evolve into the study of what amounts to be very powerful computing systems. Designing the simplest of these, i.e., the simplest universal dynamic cellular-computing system updating with natural asynchrony, is an example of a future theoretical and practical challenge for SI. More immediate, future applications can be extrapolated from the examples given throughout the entry.

Bibliography

Primary Literature

Abraham A, Grosan C, Ramos V (2006) Swarm intelligence in data mining. In: *Studies in computational intelligence*. Springer, Berlin/Heidelberg

- Adachi S, Peper F, Lee J (2004) Universality of hexagonal asynchronous totalistic cellular automata. In: 6th international conference on cellular automata for research and industry, ACRI 2004. *Lect Notes Comput Sci* 3305:91–100
- Adami C, Bryson DM, Ofria C, Pennock RT (2012) *Artificial life 13*, Ebook ISBN 9780262310505
- Agapie A, Andreica A, Chira C, Giuclea M (2014) Predictability in cellular automata. *PLoS One*. <https://doi.org/10.1371/Journal.pone.0108177>
- Aguilar W, Santamaria Bonfil G, Froese T, Gershenson C (2014) The past, present, and future of artificial life. *Front Robot AI* 1:8
- Arvin F, Murray JC, Lichen S, Chun Z, Shigang Y (2014) Development of an autonomous micro robot for swarm robotics. 2014 IEEE Int Conf Mechatron Autom 635(640):3–6
- Beni G (1988) The concept of cellular robot. In: *Proceedings of the 3rd IEEE symposium on intelligent control*, Arlington, pp 57–61
- Beni G (1992) Distributed robotic systems and swarm intelligence. *J Robot Soc Jpn* 10:31–37
- Beni G (2004a) From swarm intelligence to swarm robotics: swarm robotics. In: Sahin E, Spear WM (eds) *Revised selected papers, SAB 2004 international workshop*, Santa Monica, 17 July 2004. *Lecture notes in computer science*, March, vol 3342. Springer, pp 1–9
- Beni G (2004b) Order by disordered action in swarms. In: Sahin E, Spear WM (eds) *Revised selected papers, SAB 2004 international workshop*, Santa Monica, 17 July 17. *Lecture notes in computer science*, March, vol 3342. Springer, pp 153–171
- Beni G, Hackwood S (1992) Stationary waves in cyclic swarms. In: *Proceedings of the IEEE international symposium on intelligent control*, Glasgow, 10–13 Aug
- Beni G, Wang J (1989a) Swarm intelligence in cellular robotic systems. In: *Proceedings of NATO advanced workshop on robots and biological systems*, Tuscany, 26–30 June
- Beni G, Wang J (1989b) Swarm intelligence. In: *Proceedings for the 7th annual meeting of the robotics society of Japan*, pp 425–428 (in Japanese)
- Bersini H, Detour V (1994) Asynchrony induces stability in CA based models. In: Brooks RA, Maes P (eds) *Artificial life*, vol IV. MIT Press, Cambridge, pp 382–387
- Bonabeau E, Dorigo M, Theraulaz G (1999) *Swarm intelligence: from natural to artificial systems*. Oxford University Press, New York

References (Bonabeau et al. 1999; Kennedy et al. 2001; Abraham et al. 2006; Engelbrecht 2006; Keller et al. 2016; Dorigo and Stutzle 2004; Olariu and Zomaya 2005; Passino 2004; Gazi and Passino 2011; Yang et al. 2013) are the main books and reviews. The journal *Swarm Intelligence* (2007) is the main source for new research results. The following provide additional general material related to Swarm Intelligence.

- Brabham DC (2013) Crowdsourcing. MIT Press, Cambridge
- Brambilla M, Ferrante E, Birattari M, Dorigo M (2013) Swarm robotics: a review from the swarm engineering perspective. *Swarm Intell* 7:1–41
- Brooks R (1986) A robust layered control system for a mobile robot. *IEEE J Robot Autom* 2(1):14
- Canizo JA, Carrillo JA, Rosado J (2011) A well-posedness theory in measures for some kinetic models of collective motion. *Math Models Methods Appl Sci* 21:515–539
- Clapham N (2002) Emergent synchrony: simple asynchronous update rules can produce synchronous behavior. In: Sarker M, Gen N (eds) Proceedings of the sixth Australia-Japan joint workshop on intelligent and evolutionary systems. Australian National University, pp 41–46
- Cooper SB (2003) Computability theory. Chapman Hall/CRC, Boca Raton
- Cornforth D, Green D, Newth D (2005) Ordered asynchronous processes in multi-agent systems. *Phys D* 204(1–2):70–82
- Curtis SA, Mica J, Nuth J, Marr G, Rilee ML, Bhat M (2000) Autonomous nano-technology Swarm. In: Proceedings of the 51st international aeronautical congress, IAF-00-Q.5.08
- Dennunzio A, Formenti E, Manzoni L (2012) Computing issues of asynchronous CA. *Fundam Inform* 120:165–180
- Dorigo M (1992) Optimization, learning and natural algorithms. Ph.D thesis, Dipartimento di Elettronica, Politecnico di Milano, Milan (in Italian)
- Dorigo M, Stutzle T (2004) Ant colony optimization. MIT Press, Cambridge
- Dorigo M, Tuci E, Groß R, Trianni V, Labella TH, Nouyan S, Ampatzis C, Deneubourg J-L, Baldassarre G, Nolfi S, Mondada F, Floreano D, Gambardella LM (2004) The SWARM-BOTS project. In: Sahin E, Spears WM (eds) Proceedings of the 1st international workshop on swarm robotics. Lecture notes in computer science, vol 3342. Springer, Berlin, pp 26–40
- Dorigo M, Gambardella LM, Birattari M, Martinoli A (eds) (2006) Ant colony optimization and swarm intelligence: 5th international workshop, ANTS 2006, Brussels, 4–7 sept 2006, proceedings. Lecture notes in computer science. Springer, Berlin
- Eftimie R, de Vries G, Lewis MA (2007) Complex spatial group patterns result from different animal communication mechanisms. In: Proceedings of the National Academy of Sciences, 24 Apr 2007, vol 104, no 17
- Elamvazhuthi K, Kuiper H, Berman S (2018) PDE-based optimization for stochastic mapping and coverage strategies using robotic ensembles. *Automatica* 95:356–367. Elsevier
- Engelbrecht AP (2006) Fundamentals of computational swarm intelligence. Wiley, New York
- Fatès N (2018) Asynchronous cellular automata: a volume in the encyclopedia of complexity and systems science, 2nd ed. https://doi.org/10.1007/978-1-4939-8700-9_671
- Gardner M (1970) The fantastic combinations of John Conway's new solitaire game 'life'. *Sci Am* 223:120–123
- Gazi V, Passino KM (2011) Swarm stability and optimization. Springer, New York
- Goss S, Aron S, Deneubourg JL, Pasteel JM (1989) Self-organized shortcuts in the argentine ant. *Naturwissenschaften* 76:579–581
- Gottfredson LS (1997) Mainstream science on intelligence: an editorial with 52 signatories, history, and bibliography. *Intelligence* 24(1):13–23
- Hamann H (2010) Space-time continuous models of swarm robotic systems: supporting global-to-local programming. Springer, Berlin
- Huberman BA, Glance NS (1993) Evolutionary games and computer simulations. *Proc Natl Acad Sci U S A* 90:7716–7718
- IEEE swarm intelligence symposium. Honolulu, 1–5 Apr 2007. <http://www.computelligence.org/sis/2007/?q=node/2>
- International Journal of Swarm Intelligence Research (IJSIR) Information Resources Management Association (2010) ISSN 1947–9263
- Johnson DS, McGeoch LA (1997) The traveling salesman problem: a case study in local optimization. In: Aarts EHL, Lenstra JK (eds) Local search in combinatorial optimization. Wiley, Chichester, pp 215–310
- Johnson N, Galata A, Hogg DB (1998) The acquisition and use of interaction behavior models. In: Proceedings. 1998 IEEE computer society conference on computer vision and pattern recognition (cat. No.98CB36231), Santa Barbara, pp 866–871
- Keller JM, Liu D, Fogel DB (2016) Fundamentals of computational intelligence: neural networks, fuzzy systems, and evolutionary computation. IEEE press series on computational intelligence. Wiley/IEEE Press, Hoboken
- Kennedy J, Eberhart RC (1995) Particle swarm optimization. In: Proceedings of IEEE international conference on neural networks, vol IV. IEEE Service Center, Piscataway, pp 1942–1948
- Kennedy J, Eberhart RC, Shi Y (2001) Swarm intelligence. Morgan Kaufman, San Mateo
- Kirkpatrick S, Gelatt CD, Vecchi MP (1983) Optimization by simulated annealing. *Science* 220(4598):671–680
- Langton CG (1984) Self-reproduction in cellular automata. *Phys D* 10:135–144
- Lee J, Peper F, Adachi S, Morita K (2004) Universal delay-insensitive circuits with bi-directional and buffering lines. *IEEE Trans Comput* 53(8):1034–1046
- Legg S, Hutter M (2007) Universal Intelligence: a definition of machine intelligence. *Mind* 117(4):391–444
- Liang P, Beni G (1995) Robotic morphogenesis. *Proc Int Conf Robot Autom* 2:2175–2180
- Mandal JK, Devadutta S (2019) Intelligent computing paradigm: recent trend. Springer, Singapore
- Moere AV (2004) Information flocking: time-varying data visualization using boid behaviors. In: Proceedings of the eighth international conference on information visualization, pp 409–414

- Mogilner A, Edelstein-Keshet L, Bent L, Spiros A (2003) Mutual interactions, potentials, and individual distance in a social aggregation. *J Math Biol* 47:353–389
- Mondada F, Pettinaro GC, Guignard A, Kwee IV, Floreano D, Deneubourg J-L, Nolfi S, Gambardella LM, Dorigo M (2004) SWARM-BOT: a new distributed robotic concept. *Auton Robots* 17(2–3):193–221
- Mondada F, Gambardella LM, Floreano D, Nolfi S, Deneubourg J-L, Dorigo M (2005) The cooperation of swarm-bots: physical interactions in collective robotics. *IEEE Robot Autom Mag* 12(2):21–28
- Monmarché N (2016) Artificial ants. In: *Metaheuristics*. Springer, New York
- Murray JD (2007) *Mathematical biology I: an introduction*, 3rd edn, interdisciplinary applied mathematics. Springer, New York
- Nakamura K (1974) Asynchronous cellular automata and their computational ability. *Syst Comput Controls* 5(5):58–66
- Nehaniv CL (2002) Evolution in asynchronous cellular automata. In: Standish RK, Abbass HA, Bedau MA (eds) *Proceedings of the eighth conference on artificial life*. MIT Press, pp 65–74
- Nicolis G, Prigogine I (1977) *Self-organization in non-equilibrium systems*. Wiley, New York
- Nor Azlina AA, Mubin M, Mohamad MS, Kamarulzaman AA (2014) A synchronous-asynchronous particle swarm optimization algorithm. *Sci World J* 2014:123019. <https://doi.org/10.1155/2014/123019>
- Olariu S, Zomaya AY (2005) *Handbook of bioinspired algorithms and applications*. Chapman & Hall/CRC Computer & Information Science, Boca Raton
- Parker LE, Schneider FE, Schultz AC (2005) Multi-robot systems. From swarms to intelligent automata. In: *Proceedings from the 2005 international workshop on multi-robot systems*, vol III. Springer
- Passino K (2004) *Biomimicry for optimization, control, and automation*. Springer, London
- Peruggia M (1993) *Discrete iterated function systems*. CRC Press, Wellesley. 1568810156
- Pincirolì C, Trianni V, O’Grady R, Pini G, Brutschy A, Brambilla M, Mathews N, Ferrante E, DiCaro G, Ducatelle F, Birattari M, Gambardella LM, Dorigo M (2012) ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems. *Swarm Intell* 6:271–295
- Reynolds C (1987) Flocks, herds, and schools: a distributed behavioral model. *Comput Graph* 21(4):25–34
- Rubenstein M, Comejo A, Nagpal R (2014) Programmable self-assembly in a thousand-robot swarm. *Science* 345:6198
- Sahin E, Spears WM (2005) *Swarm robotics: SAB 2004 international workshop*, Santa Monica, 17 July 2004, revised selected papers. Lecture notes in computer science. Springer
- Sahin E, Spears WM, Winfield AFT (eds) (2007) *Swarm robotics*. Second SAB 2006 international workshop, Rome, 30 Sept 2006–1 Oct 2006 revised selected papers. Lecture notes in computer science, vol 4433. Springer, Berlin/Heidelberg/New York
- Schonfisch B, de Roos A (1999) Synchronous and asynchronous updating in cellular automata. *Biosystems* 51:123–143
- Sengupta S, Basak S, Peters II RA (2018) Particle swarm optimization: a survey of historical and recent developments with hybridization perspectives. *Mach Learn Knowl Extr* 1:157–191
- Sipper M (1997) *Evolution of parallel cellular machines: the cellular programming approach*. Lecture notes in computer science. Springer, New York
- Sipper M (1999) The emergence of cellular computing. *IEEE Comput* 32(7):18–26
- Sipper M, Tomassini M, Capcarrere MS (1997) Evolving asynchronous and scalable non-uniform cellular automata. In: *Proceedings of international conference on artificial neural networks and genetic algorithms (ICANNGA97)*. Springer
- Snyman JA (2005) *Practical mathematical optimization*. An introduction to basic optimization theory and classical and new gradient-based algorithms. Springer, New York
- Sun J, Lai C-H, Wu X-J (2011) Particle swarm optimisation: classical and quantum perspectives. Numerical analysis and scientific computing series. Chapman & Hall/CRC, 1 Har/Cdr, Boca Raton
- Swarm Intelligence (2007) Springer. ISSN: 1935–3812
- Takada Y, Isokawa T, Peper F, Matsui N (2006) Construction universality in purely asynchronous cellular automata. *J Comput Syst Sci* 72:1368–1385
- Tan Y (ed) (2018) *Swarm Intelligence: principles, current algorithms and methods (control, robotics, and sensors)*. The Institution of Engineering and Technology, London
- Topaz CM, Bertozzi A (2004) Swarming patterns in two-dimensional kinematic model for biological groups. *SIAM J Appl Math* 65(1):152–174
- Topaz CM, D’Orsogna MR, Edelstein-Keshet L, Bernoff AJ (2012) Locust dynamics: behavioral phase change and swarming. *PLOS Comp Biol* 8(8):e1002642
- Traub J (editor in chief) *J Complex*. Elsevier http://www.elsevier.com/wps/find/journaldescription.cws_home/622865/description#description
- Trianni V, Tuci E, Passino KM, Marshall JAR (2011) *Swarm cognition: an interdisciplinary approach to the study of self-organising biological collectives*. *Swarm Intell* 5:3–18
- Turing AM (1952) The chemical basis for morphogenesis. *Philos Trans R Soc Lond B* 237:37–72
- von Neumann J (1966) *Theory of self-reproducing automata*. University of Illinois Press. edited and completed by Burks AW
- Weiss G (2000) *Multiagent systems: a modern approach to distributed artificial intelligence*. MIT Press, Cambridge
- Wolfram S (1985) Undecidability and intractability in theoretical physics. *Phys Rev Lett* 54:735–738
- Wolfram S (2002) *A new kind of science*. Wolfram Media, Champaign
- Xu W-L, Lee J, Chen H-H, Isokawa T (2019) Universal computation in a simplified Brownian cellular automaton with Von Neumann Neighborhood. *Fundamenta Informaticae* 165(2):139–156

Yang X-S, Cui Z, Xiao R, Gandomi AH (2013) *Swarm intelligence and bio-inspired computation: theory and applications*. Elsevier, Boston

Books and Reviews

- Agrawal A, Gans J, Goldfarb A (2018) *Prediction machines: the simple economics of artificial Intelligence*. Harvard Business Review Press, Boston
- Camazine S, Deneubourg J-L, Franks NR, Sneyd J, Theraulaz G, Bonabeau E (2001) *Self-organization in biological systems*. Princeton University Press, Princeton
- Deutsch A, Dormann S (2018) *Cellular automaton modeling of biological pattern formation: characterization, examples, and analysis*, 2nd edn. Birkhauser, Basel
- Dorigo M, Sahin E (2004) *Swarm robotics – special issue editorial*. *Auton Robot* 17(2–3):111–113
- Engelbrecht AP (2006) *Fundamentals of computational swarm intelligence*. Wiley, New York
- Hamann H (2010) *Space-time continuous models of swarm robotic systems: supporting global-to-local programming, cognitive systems monographs*. Springer, Berlin
- Hamann H (2018) *Swarm robotics: a formal approach*. Springer, Cham
- Hassanien AE, Emary E (2016) *Swarm intelligence: principles, advances, and applications*. CRC Press, Boca Raton
- Kacprzyk J, Pedrycz W (2015) *Springer handbook of computational intelligence*. Springer, Berlin
- Kruse R, Borgelt C, Klawonn F, Moewes C, Steinbrecher M, Held P (2013) *Computational intelligence: a methodological introduction*. Springer, New York
- Mohanty S (2018) *Swarm Intelligence methods for statistical regression*. Chapman and Hall/CRC, Boca Raton
- Nayyar A, Le D-N, Nguyen NG (2018) *Advances in swarm intelligence for optimizing problems in computer science*. Chapman and Hall/CRC, Boca Raton
- Sipper M (2002) *Machine nature: the coming age of bio-inspired computing*. McGraw-Hill, New York
- Solnon C (2010) *Ant colony optimization and constraint programming*. Wiley-ISTE, Hoboken
- Tan Y, Takagi H, Shi Y, Niu B (2017) *Advances in swarm intelligence: 8th international conference, ICSI 2017, Fukuoka*

Social Phenomena Simulation

Paul Davidsson¹ and Harko Verhagen²

¹Department of Computer Science, Malmö University, Malmö, Sweden

²Department of Computer and Systems Sciences, Stockholm University, Kista, Sweden

Glossary

Agent (or software agent) A self-contained entity that has a state and that is situated (able to perceive and act) in an environment. In addition, agents are often assumed to be rational and autonomous.

Cellular automaton A mathematical structure modeling a set of cells that interact with their neighbors. Each cell has a set of neighbors and a state. All the cells update their values simultaneously at discrete time steps. The new state of a cell is determined by the current state of its neighbors according to a local function or rule.

Microlevel simulation A type of simulation in which the specific behaviors of specific individuals are explicitly modeled.

Definition of the Subject

Social phenomena simulation in the area of agent-based modeling and simulation concerns the emulation of the individual behavior of a group of social entities, typically including their cognition, actions, and interaction. Agent-based social simulation constitutes the intersection of three scientific fields, namely, agent-based computing, the social sciences, and computer simulation (Davidsson 2000). Agent-based computing is a research area mainly within computer science and includes, e.g., agent-based modeling, design, and programming. By the social sciences, we here refer to a large set of different sciences that study the interaction among social entities, e.g., social

psychology, management science, policy, and some areas of biology. Computer simulation concerns the study of different techniques for simulating phenomena on a computer, e.g., discrete-event, object-oriented, and equation-based simulation.

Introduction

Computer simulation consists of three main steps: (i) designing a model of an actual or theoretical system, (ii) executing the model on a computer, and (iii) analyzing the execution output. Already in the early days of computer development, simulation was used in different research areas to predict the behavior of complex systems. Such simulations were typically based on differential equations and focused on results at the aggregate level. These models of, for instance, predator-prey populations could result in fairly accurate models but were limited in the sense that the models excluded individual behavior and decision making, as well as interaction between individuals, and were based on homogeneous agents. The development of agent-based modeling offers a possible solution to this problem with its (seemingly) natural mapping onto interacting individuals with incomplete information and capabilities, no global control, decentralized data, asynchronous computing, and inclusion of heterogeneous agents. Agent-based simulation models also offer the possibility of studying the dynamics of the interaction processes instead of focusing on the (static) results of these processes (Sawyer 2003; Mustaphaa et al. 2013).

Agent-based modeling can be traced back to von Neumann, who in the 1950s invented what was later termed *cellular automata*. These were used by Conway in the 1970s when he constructed the well-known *Game of Life*. It is based on very simple rules determining the life and death of the cells in a virtual world in the form of a 2D grid. Inspired by this work, researchers developed more-refined models, often modeling

the social behavior of groups of animals or artificial creatures. One example is the Boid model by Reynolds (1987), which simulates coordinated animal motion such as bird flocks and fish schools. With respect to human societies, Epstein and Axtell (1996) developed one of the first agent-based models, called Sugarscape, to explore the role of social phenomena such as seasonal migrations, pollution, sexual reproduction, combat, and transmission of disease. This work is in spirit closely related to one of the best-known and earliest examples of the use of simulation in social science, namely, the Schelling model (1971), in which cellular automata were used to simulate the emergence of segregation patterns in neighborhoods based on a few simple rules expressing the preferences of the agents. Another pioneer worth mentioning is Barricelli (1957), who to some extent used agent-based modeling for simulating biological systems.

To sum up, we can identify two main approaches to social simulation:

- Macrolevel (or equation-based) simulation, which is typically based on mathematical models. It views the set of individuals (the population) as a structure that can be characterized by a number of variables.
- Microlevel (or agent-based) simulation, in which the specific behaviors of specific individuals are explicitly modeled. In contrast to macrolevel simulation, it views the structure as emerging from the interactions between individuals and thus exploring the standpoint that complex effects need not have complex causes.

As argued by Van Parunak et al. (1998), agent-based modeling is most appropriate for domains characterized by a high degree of localization and distribution and dominated by discrete decision. Equation-based modeling, on the other hand, is most naturally applied to systems that can be modeled centrally and in which the dynamics are dominated by physical laws rather than information processing. We will here focus on agent-based models, particularly those that have a richer representation of the individual than the cellular automata.

Why Simulate Social Phenomena?

Simulation of social phenomena can be done for different purposes, e.g.,

- Supporting social-theory building
- Supporting the engineering of systems, e.g., validation, testing, etc.
- Supporting planning, policymaking, and other decision making
- Training, in order to improve a person's skills in a certain domain

It is possible to distinguish between four types of end users: *scientists*, who use social phenomena simulation in the research process to gain new knowledge; *policymakers*, who use it for making strategic decisions; *managers* (of systems), who use it to make operational decisions; and *other professionals*, such as architects, who use it in their daily work. We will now describe how these types of end users may use simulation of social phenomena for different purposes.

Supporting Social-Theory Building

In the context of social-theory building, agent-based simulation can be seen as an experimental method or as theories in themselves (Sawyer 2003). In the former case, simulations are run to test the predictions of theories, whereas in the latter case, simulations in themselves are formal models of theories. Formalizing the ambiguous, natural language-based theories of the social sciences helps to find inconsistencies and other problems and thus contributes to theory building.

Using agent-based simulation studies as an experimental tool offers great possibilities. Many experiments with human societies are either unethical or even impossible to conduct. Experiments *in silico*, on the other hand, are fully possible. These can also breathe new life into the ever-present debate in sociology on the micro-macro link (Alexander et al. 1987). Agent-based models mostly focus on the emergence of macrolevel properties from the local interaction of adaptive agents that influence one another (Macy and

Miller 2002; Sawyer 2003). However, simulations in computational organization theory (Carley and Prietula 1994; Prietula et al. 1998), for example, often try to analyze the influence of macrolevel phenomena on individuals. Using agent-based models to simulate the bidirectional relation between micro- and macrolevel concepts would provide tools to analyze the theoretical consequences of the work done by theorists such as Habermas, Giddens, and Bourdieu, to name a few (Sawyer 2003).

Supporting the Engineering of Systems

Many new technical systems are distributed and involve complex interactions between humans and machines. The properties of agent-based simulation make it especially suitable for simulating these kinds of systems. The idea is to model the behavior of human users in terms of software agents. In particular, this seems useful in situations where it is too expensive, difficult, inconvenient, tiresome, or even impossible for real human users to test out a new technical system. Of course, also the technical system, or parts thereof, may be simulated. For instance, if the technical system includes hardware that is expensive and/or special purpose, it is natural to simulate also this part of the system when testing out the control software. An example of such a case is the testing of control systems for “intelligent buildings,” where agents simulate the behavior of the people in the building (Davidsson 2000).

Supporting Planning, Policymaking, and Other Decision Making

Here the focus is on exploring different possible future scenarios in order to choose between alternative actions. Besides this type of prediction, simulation of social phenomena may be used for analysis, i.e., to gain deeper knowledge and understanding of a certain phenomenon.

An area in which several studies of this kind have been carried out is disaster management, such as experiments concerning different roles

and the efficiency of reactions to emergencies (Friedricj and Burghardt 2007). For a recent overview of work in this area, see Mustaphaa et al. (2013). An example of the use of real-world data is to analyze the effect of different insurance policies on the willingness of agents to pay for a disaster insurance policy (Brouwers and Verhagen 2004). In the related area of flood management, Dawson et al. (2011) describe a model where individual decisions by inhabitants are implemented as probabilistic finite state machines where the parameters are linked to census data.

Another application area for this type of simulation study is disease spreading. Typically, agents are used to represent human beings and the simulation model is linked to real-world geographical data. One study (Yergens et al. 2006) also included agents that represent towns acting as the epicenter of disease outbreak. The town agent’s behavior repertoire consisted of different containment strategies. The simulation model can be quickly adapted to local circumstances via the geographical data (given that there is data on the population as well) and is used to determine the effects of different containment strategies.

A third area where agent-based social simulation has been used to support planning and policymaking is traffic and transport (Bazzan and Kügl 2014). An example of this is the simulation of all cars traveling in Switzerland during morning peak traffic (Balmer et al. 2008).

Training

The main advantage of using simulation for training purposes is to be part of a real-world-like situation without real-world consequences. Especially in the military, the use of simulation for training purposes is widespread. Also in medicine, where mistakes can be very expensive in terms of money and lives, the use of simulation in education is on the rise.

An early product in this area was a tool to help train police officers to manage large public gatherings such as crowds, demonstrations, and marches (Williams 1993). Another early example of agent-based simulation for training purposes is

Steve (Rickel and Lewis Johnson 1999, Méndez et al. 2003). Steve was an agent integrated with voice synthesis software and virtual reality software providing a very realistic training environment for controlling the engine room of a virtual US Navy surface ship.

Another example is the PSI agent (Künzel and Hammer 2006). Whereas in most cases the simulator training is aimed at training practical skills or decision making, this work focuses on acquiring theoretical insights in the realm of psychological theory. The simulation enables students to explore psychological processes without ethical problems.

Simulating Social Phenomena

One of the first, and simplest, ways of performing microlevel simulation is often called *dynamic microsimulation* (Gilbert 1999; Gilbert and Troitzch 2005). It is used to simulate the effect of the passing of time on individuals. Data from a (preferably large) random sample from the population to be simulated is used to initially characterize the simulated individuals. Some examples of sampled features are age, sex, employment status, income, and health status. A set of transition probabilities are used to describe how these features will change over a given time period, e.g., there is a probability that an employed person will become unemployed over the course of a year. The transition probabilities are applied to the population for each individual in turn and then repeatedly reapplied for a number of simulated time periods. Sometimes it is necessary to also model changes in the population, e.g., birth, death, and marriage. This type of simulation can be used to, e.g., predict the outcome of different social policies. However, the quality of such simulations depends on the quality of the following:

- The random sample, which must be representative
- The transition probabilities, which must be valid and complete

In traditional microsimulation, the behavior of each individual is regarded as a “black box.”

The behavior is modeled in terms of probabilities, and no attempt is made to justify these in terms of individual preferences, decisions, plans, etc. Also, each simulated individual is considered in isolation without regard to their interaction with others. Thus, better results may be gained if cognitive processes and communication between individuals are also simulated.

Opening the black box of individual decision making can be done in several ways. The first layer to add is often individual psychology; for instance, the so-called beliefs, desires, and intentions (BDI) model is often used. For an overview of simulations using BDI models, see Adam and Gadou (2016). Models of individual cognition used in agent-based social simulation include the use of Soar (a computer implementation of Allen Newell’s unified theory of cognition (Künzel and Hämmer 2006)), which was used in Steve (discussed in section “Why Simulate Social Phenomena?”).

For the simulation of social behavior, the agents need to be equipped with mechanisms for reasoning at the social level (unless the social level is regarded as emerging from individual behavior and decision making). Several models have been based on theories from economics, social psychology, sociology, etc. An early example of this is provided by Guye-Vuillème (2004), who has developed an agent-based model for simulating human interaction in a virtual reality environment. The model is based on sociological concepts such as roles, values, and norms and motivational theories from social psychology to simulate persons with social identities and relationships. Another example is the Consumat model (Janssen and Jager 1999), a metamodel combining several psychological theories on decision making in a consumer situation, used, for instance, to investigate different flood management policies (Brouwers and Verhagen 2004).

Future Directions

In a study of applications of agent-based simulation (Davidsson et al. 2007), it was concluded that even if agent-based simulation seems a promising

approach to many problems involving the simulation of complex systems of interacting entities such as social phenomena, it seems that the full potential of the agent concept often is not utilized. For instance, most models have very primitive agent cognition, in particular if the number of agents involved is large.

Regarding future applications, the combination of findings from neuroscience and simulation studies as proposed by Epstein (2014) is a new development and an interesting follow-up to the abovementioned Sugarscape book. The wide availability of high-quality graphics opens for agent-based social simulation in combination with sophisticated visualization techniques, such as virtual reality, in the form of “serious games,” which has the potential to provide very powerful training environments; see, e.g., Garcia Carbajal et al. (2015). In the context of military training, McAlinden et al. (2014) provide an interesting application.

Bibliography

- Adam C, Gaudou B (2016) BDI agents in social simulations: a survey. *Knowl Eng Rev* 31(3):207–238
- Alexander JC, Giesen B, Münch R, Smelser NJ (eds) (1987) *The micro-macro link*. University of California Press, Berkeley
- Balmer M, Meister K, Nagel K (2008) Agent-based simulation of travel demand: structure and computational performance of MATSim-T. ETH, Eidgenössische Technische Hochschule Zürich, IVT Institut für Verkehrsplanung und Transportsysteme
- Barricelli NA (1957) Symbiogenetic evolution processes realized by artificial methods. *Methodos* 9(35–36):143–182
- Bazzan A, Klügl F (2014) A review on agent-based technology for traffic and transportation. *Knowl Eng Rev* 29(3):375–403
- Brouwers L, Verhagen H (2004) Applying the Consumat model to flood management policies. In: Müller J, Seidel M (eds) 4th workshop on agent-based simulation. SCS, Montpellier, pp 29–34
- Carley KM, Prietula M (eds) (1994) *Computational organization theory*. Erlbaum, Hillsdale
- Davidsson P (2000) Multi-agent-based simulation: beyond social simulation. In: Moss S, Davidsson P (eds) *Multi-agent-based simulation, Lecture notes in computer science*, vol 1979. Springer, Berlin, pp 98–107
- Davidsson P (2002) Agent-based social simulation: a computer science view. *J Artif Soc Soc Simul* 5(1)
- Davidsson P, Holmgren J, Kyhlbäck H, Mengistu D, Persson M (2007) Applications of multi-agent-based simulation. In: Antunes L, Takadama K (eds) *Multi-agent-based simulation VII, Lecture notes in computer science*, vol 4442. Springer, Berlin
- Dawson RJ, Peppe R, Wang M (2011) An agent-based model for risk-based flood incident management. *Nat Hazards* 59(1):167–189
- Epstein JM (2014) *Agent_zero: toward neurocognitive foundations for generative social science*. Princeton University Press
- Epstein JM, Axtell RL (1996) *Growing artificial societies: social science from the bottom up*. MIT Press, Cambridge
- Fiedrich F, Burghardt P (2007) Emergency response information systems: emerging trends and technologies: agent-based systems for disaster management. *Commun ACM* 50(3):41–42
- García Carbajal S, Polimeni F, Múgica JL (2015) An emotional engine for behavior simulators. *Int J Serious Game* 2(2):57–67
- Gilbert N (1999) *Computer simulation in the social sciences*. Sage, Thousand Oaks
- Gilbert N, Troitzsch KG (2005) *Simulation for the social scientist*, 2nd edn. Open University Press, Maidenhead
- Guye-Vuillème A (2004) *Simulation of nonverbal social interaction and small groups dynamics in virtual environments*. Ph D thesis, École Polytechnique Fédérale de Lausanne, No 2933
- Janssen MA, Jager W (1999) An integrated approach to simulating behavioural processes: a case study of the lock-in of consumption patterns. *J Artif Soc Soc Simul* 2(2). <http://jasss.soc.surrey.ac.uk/2/2/2.html>
- Künzel J, Hämmer V (2006) Simulation in university education: the artificial agent PSI as a teaching tool. *Simulation* 82(11):761–768
- Macy MW, Willer R (2002) From factors to actors: computational sociology and agent-based modeling. *Annu Rev Sociol* 28:143–166
- McAlinden R, Pynadath D, Hill RW (2014) UrbanSim: using social simulation to train for stability operations. In: Ehlschlaeger CH (ed) *Understanding megacities with the reconnaissance, surveillance, and intelligence paradigm*, US Army white report
- Méndez G, Rickel J, de Antonio A (2003) Steve meets Jack: the integration of an intelligent tutor and a virtual environment with planning capabilities. In: *Intelligent virtual agents, Lecture notes on artificial intelligence*, vol 2792. Springer, Berlin, pp 325–332
- Mustaphaa K, Mcheicka H, Melloulis S (2013) Modeling and simulation agent-based of natural disaster complex systems. *Procedia Comput Sci* 21:148–155. <https://doi.org/10.1016/j.procs.2013.09.021>
- Newell A (1994) *Unified theories of cognition*. Harvard University Press, Cambridge
- Parunak HVD, Savit R, Riolo RL (1998) Agent-based modeling vs. equation-based modeling: a case study and users’ guide. In: Sichman JS, Conte R, Gilbert N (eds) *Multi-agent systems and agent-based simulation, Lecture notes in computer science*, vol 1534. Springer, Berlin, pp 10–26
- Prietula MJ, Carley KM, Gasser L (eds) (1998) *Simulating organizations: computational models of institutions and groups*. MIT Press, Cambridge

- Reynolds CW (1987) Flocks, herds, and schools: a distributed behavioral model. *Comput Graph* 21(4):25–34. (SIGGRAPH'87 conference proceedings)
- Rickel J, Lewis Johnson W (1999) Animated agents for procedural training in virtual reality: perception, cognition, and motor control. *Appl Artif Intell* 13:343–382
- Sawyer RK (2003) Artificial societies – multi-agent systems and the micro-macro link in sociological theory. *Sociol Methods Res* 31(3):325–363
- Schelling TC (1971) Dynamic models of segregation. *J Math Sociol* 1(1):143–186
- Takahashi S, Sallach D, Rouchier J (2007) Advancing social simulation – the first world congress. Springer, Berlin
- Williams R (1993) An agent based simulation environment for public order management training. In: Western simulation multiconference, object-oriented simulation conference (OOS '93), San Diego, pp 151–156
- Yergens D, Hiner J, Denzinger J, Noseworthy T (2006) IDESS – a multi-agent-based simulation system for rapid development of infectious disease models. *Int Trans Syst Sci Appl* 1(1):51–58

Further Reading

Classic books in the area of simulation of social behavior include “Growing Artificial Societies: Social Science from the Bottom Up” by Epstein and Axtell (1996) and “Simulation for the Social Scientist” by Gilbert and Troitzsch (2005)

More recent findings can be found in, e. g., the *Journal of Artificial Societies and Social Simulation* (<http://jasss.soc.surrey.ac.uk>) and the proceedings of, e. g., the International Workshop series on Multi-Agent-Based Simulation (MABS) (<http://www.pcs.usp.br/~mabs/>), the series of International Conference on Social Computing, Behavioral-Cultural Modeling, & Prediction and Behavior Representation in Modeling and Simulation (<http://sbp-brims.org>) and the conference of the European Social Simulation Association (ESSA) (<http://www.essa.eu.org/>), the Pacific-Asian Association for Agent-based Approach in Social Systems Sciences (<http://www.paaa.asia/>) and the Computational Social Science Society of the Americas (<http://www.computationalsocialscience.org>)

Agent-Based Computational Economics

Moshe Levy
The Hebrew University, Jerusalem, Israel

Article Outline

Glossary
Definition of the Subject
Introduction
Some of the Pioneering Studies
Illustration with the LLS Model
Summary and Future Directions
Bibliography

Keywords

microscopic simulation · agent-based simulation · financial markets · bounded rationality · heterogeneous expectations

Glossary

Agent-based simulation A simulation of a system of multiple interacting agents (sometimes also known as “microscopic simulation”). The “micro” rules governing the actions of the agents are known and so are their rules of interaction. Starting with some initial conditions, the dynamics of the system are investigated by simulating the state of the system through discrete time steps. This approach can be employed to study general properties of the system, which are not sensitive to the initial conditions, or the dynamics of a specific system with fairly well-known initial conditions, e.g., the impact of the baby boomers’ retirement on the US stock market.

Bounded rationality Most economic models describe agents as being *fully* rational – given the information at their disposal, they act in the

optimal way which maximizes their objective (or utility) function. This optimization may be technically very complicated requiring economic, mathematical, and statistical sophistication. In contrast, *bounded*-rational agents are limited in their ability to optimize. This limitation may be due to limited computational power, errors, or various psychological biases which have been experimentally documented.

Market anomalies Empirically documented phenomena that are difficult to explain within the standard rational-representative-agent economic framework. Some of these phenomena are the overreaction and underreaction of prices to news the autocorrelation of stock returns, various calendar and day-of-the-week effects, and the excess volatility of stock returns.

Representative agent A standard modeling technique in economics by which an entire class of agents (e.g., investors) are modeled by a single “representative” agent. If agents are completely homogeneous, it is obvious that the representative agent method is perfectly legitimate. However, when agents are heterogeneous, the representative agent approach can lead to a multitude of problems (see Kirman 1992).

Definition of the Subject

Mainstream economic models typically make the assumption that an entire group of agents, e.g., “investors,” can be modeled with a single “rational representative agent.” While this assumption has proven extremely useful in advancing the science of economics by yielding analytically tractable models, it is clear that the assumption is not realistic: people are different one from the other in their tastes, beliefs, and sophistication, and as many psychological studies have shown, they often deviate from rationality in systematic ways.

Agent-based computational economics is a framework allowing economics to expand beyond the realm of the “rational representative agent.” Modeling and simulating the behavior of each

agent and the interaction among agents agent-based simulation allows us to investigate the dynamics of complex economic systems with many heterogeneous and not necessarily fully rational agents.

The agent-based simulation approach allows economists to investigate systems that cannot be studied with the conventional methods. Thus, the following key questions can be addressed: How do heterogeneity and systematic deviations from rationality affect markets? Can these elements explain empirically observed phenomena which are considered “anomalies” in the standard economics literature? How robust are the results obtained with the analytical models? By addressing these questions the agent-based simulation approach complements the traditional analytical analysis and is gradually becoming a standard tool in economic analysis.

Introduction

For solving the dynamics of two bodies (e.g., stars), with some initial locations and velocities and some law of attraction (e.g., gravitation), there is a well-known analytical solution. However, for a similar system with three bodies, there is no known analytical solution. Of course, this does not mean that physicists cannot investigate and predict the behavior of such systems. Knowing the state of the system (i.e., the location, velocity, and acceleration of each body) at time t allows us to calculate the state of the system an instant later, at time $t + \Delta t$. Thus, starting with the initial conditions, we can predict the dynamics of the system by simply simulating the “behavior” of each element in the system over time.

This powerful and fruitful approach, sometimes called “microscopic simulation,” has been adopted by many other branches of science. Its application in economics is best known as “agent-based simulation” or “agent-based computation.” The advantages of this approach are clear – they allow the researcher to go where no analytical models can go. Yet, despite of the advantages, perhaps surprisingly, the agent-based approach was not adopted very quickly by economists. Perhaps the main reason for this is that a particular simulation only describes the dynamics of a system with a particular set of parameters and initial conditions. With other

parameters and initial conditions, the dynamics may be different. So economists may ask: what is the value of conducting simulations if we get very different results with different parameter values? While in physics the parameters (like the gravitational constant) may be known with great accuracy, in economics the parameters (like the risk-aversion coefficient or for that matter the entire decision-making rule) are typically estimated with substantial error. This is a strong point. Indeed, we would argue that the “art” of agent-based simulations is the ability to understand the general dynamics of the system and to draw general conclusions from a finite number of simulations. Of course, one simulation is sufficient as a counterexample to show that a certain result does not hold, but many more simulations are required in order to convince of an alternative general regularity.

This manuscript is intended as an introduction to agent-based computational economics. An introduction to this field has two goals: (i) to explain and to demonstrate the agent-based methodology in economics, stressing the advantages and disadvantages of this approach relative to the alternative purely analytical methodology, and (ii) to review studies published in this area. The emphasis in this entry will be on the first goal. While section “[Some of the Pioneering Studies](#)” does provide a brief review of some of the cornerstone studies in this area, more comprehensive reviews can be found in LeBaron (2000), Levy et al. (2000), Samanidou et al. (2007), and Tesfatsion (2001, 2002), in which part of section “[Some of the Pioneering Studies](#)” is based. A comprehensive review of the many entries employing agent-based computational models in economics will go far beyond the scope of this entry. To achieve goal (i) above, in section “[Illustration with the LLS Model](#),” we will focus on one particular model of the stock market in some detail. Section “[Summary and Future Directions](#)” concludes with some thoughts about the future of the field.

Some of the Pioneering Studies

Schelling’s Segregation Model

Schelling’s (Schelling 1978) classical segregation model is one of the earliest models of population

dynamics. Schelling's model is not intended as a realistic tool for studying the actual dynamics of specific communities as it ignores economic, real-estate, and cultural factors. Rather, the aim of this very simplified model is to explain the emergence of macroscopic single-race neighborhoods even when individuals are not racists. More precisely, Schelling found that the collective effect of neighborhood racial segregation results even from individual behavior that presents only a very mild preference for same-color neighbors. For instance, even the minimal requirement by each individual of having (at least) one neighbor belonging to one's own race leads to the segregation effect.

The agent-based simulation starts with a square mesh or lattice (representing a town) which is composed of cells (representing houses). On these cells reside agents which are either "blue" or "green" (the different races). The crucial parameter is the minimal percentage of same-color neighbors that each agent requires. Each agent, in his/her turn, examines the color of all his/her neighbors. If the percentage of neighbors belonging to his/her own group is above the "minimal percentage," the agent does nothing. If the percentage of neighbors of his/her own color is less than the minimal percentage, the agent moves to the closest unoccupied cell. The agent then examines the color of the neighbors of the new location and acts accordingly (moves if the number of neighbors of his/her own color is below the minimal percentage and stays there otherwise). This goes on until the agent is finally located at a site in which the minimal percentage condition holds. After a while, however, it might happen that following the moves of the other agents, the minimal percentage condition ceases to be fulfilled and then the agent starts moving again until he/she finds an appropriate cell. As mentioned above, the main result is that even for very mild individual preferences for same-color neighbors, after some time the entire system displays a very high level of segregation.

A more modern, developed, and sophisticated reincarnation of these ideas is the Sugarscape environment described by Epstein and Axtell (1996). The model considers a population of moving, feeding, pairing, procreating, trading, warring agents and displays various qualitative collective events which their populations incur. By

employing agent-based simulation, one can study the macroscopic results induced by the agents' individual behavior.

The Kim and Markowitz Portfolio Insurers Model

Harry Markowitz is very well known for being one of the founders of modern portfolio theory, a contribution for which he/she has received the Nobel Prize in economics. It is less well known, however, that Markowitz is also one of the pioneers in employing agent-based simulations in economics.

During the October 1987 crash, markets all over the globe plummeted by more than 20% within a few days. The surprising fact about this crash is that it appeared to be spontaneous – it was not triggered by any obvious event. Following the 1987 crash, researchers started to look for endogenous market features, rather than external forces, as sources of price variation. The Kim-Markowitz (Kim and Markowitz 1989) model explains the 1987 crash as resulting from investors' "Constant Proportion Portfolio Insurance" (CPPI) policy. Kim and Markowitz proposed that market instabilities arise as a consequence of the individual insurers' efforts to cut their losses by selling once the stock prices are going down.

The Kim-Markowitz agent-based model involves two groups of individual investors: rebalancers and insurers (CPPI investors). The rebalancers are aiming to keep a constant composition of their portfolio, while the insurers make the appropriate operations to insure that their eventual losses will not exceed a certain fraction of the investment per time period.

The rebalancers act to keep a portfolio structure with (for instance) half of their wealth in cash and half in stocks. If the stock price rises, then the stocks weight in the portfolio will increase, and the rebalancers will sell shares until the shares again constitute 50% of the portfolio. If the stock price decreases, then the value of the shares in the portfolio decreases, and the rebalancers will buy shares until the stock again constitutes 50% of the portfolio. Thus, the rebalancers have a stabilizing influence on the market by selling when the market rises and buying when the market falls.

A typical CPPI investor has as his/her main objective not to lose more than (for instance) 25 % of his/her initial wealth during a quarter, which consists of 65 trading days. Thus, he/she aims to insure that at each cycle 75% of the initial wealth is out of reasonable risk. To this effect, he/she assumes that the current value of the stock will not fall in one day by more than a factor of 2. The result is that he/she always keeps in stock twice the difference between the present wealth and 75% of the initial wealth (which he/she had at the beginning of the 65-day investing period). This determines the amount the CPPI agent is bidding or offering at each stage. Obviously, after a price fall, the amount he/she wants to keep in stocks will fall and the CPPI investor will sell and further destabilize the market. After an increase in the prices (and personal wealth), the amount the CPPI agent wants to keep in shares will increase: he/she will buy and may support a price bubble.

The simulations reveal that even a relatively small fraction of CPPI investors (i.e., less than 50%) is enough to destabilize the market, and crashes and booms are observed. Hence, the claim of Kim and Markowitz that the CPPI policy may be responsible for the 1987 crash is supported by the agent-based simulations. Various variants of this model were studied intensively by Egener et al. (1999) who find that the price time evolution becomes unrealistically periodic for a large number of investors (the periodicity seems related with the fixed 65-day quarter and is significantly diminished if the 65-day period begins on a different date for each investor).

The Arthur, Holland, Lebaron, Palmer, and Tayler Stock Market Model

Palmer et al. (1994) and Arthur et al. (1997) (AHLPT) construct an agent-based simulation model that is focused on the concept of coevolution. Each investor adapts his/her investment strategy such as to maximally exploit the market dynamics generated by the investment strategies of all others investors. This leads to an ever-evolving market, driven endogenously by the ever-changing strategies of the investors.

The main objective of AHLPT is to prove that market fluctuations may be induced by this

endogenous coevolution, rather than by exogenous events. Moreover, AHLPT study the various regimes of the system: the regime in which rational fundamentalist strategies are dominating versus the regime in which investors start developing strategies based on technical trading. In the technical trading regime, if some of the investors follow fundamentalist strategies, they may be punished rather than rewarded by the market. AHLPT also study the relation between the various strategies (fundamentals vs. technical) and the volatility properties of the market (clustering, excess volatility, volume-volatility correlations, etc.).

In the first entry quoted above, the authors simulated a single stock and further limited the bid/offer decision to a ternary choice of: (i) bid to buy one share, (ii) offer to sell one share, or: (iii) do nothing. Each agent had a collection of rules which described how he/she should behave (i, ii, or iii) in various market conditions. If the current market conditions were not covered by any of the rules, the default was to do nothing. If more than one rule applied in a certain market condition, the rule to act upon was chosen probabilistically according to the “strengths” of the applicable rules. The “strength” of each rule was determined according to the rule’s past performance: rules that “worked” became “stronger.” Thus, if a certain rule performed well, it became more likely to be used again.

The price is updated proportionally to the relative excess of offers over demands. In Arthur et al. (1997), the rules were used to predict future prices. The price prediction was then transformed into a buy/sell order through the use of a Constant Absolute Risk-Aversion (CARA) utility function. The use of CARA utility leads to demands which do not depend on the investor’s wealth.

The heart of the AHLPT dynamics is the trading rules. In particular, the authors differentiate between “fundamental” rules and “technical” rules and study their relative strength in various market regimes. For instance, a “fundamental” rule may require market conditions of the type:

$$\text{dividend/current price} > 0.04$$

in order to be applied. A “technical” rule may be triggered if the market fulfills a condition of the type:

current price $>$
 10 – period moving average of past prices

The rules undergo genetic dynamics: the weakest rules are substituted periodically by copies of the strongest rules, and all the rules undergo random mutations (or even versions of “sexual” crossovers – new rules are formed by combining parts from two different rules). The genetic dynamics of the trading rules represent investors’ learning: new rules represent new trading strategies. Investors examine new strategies and adopt those which tend to work best. The main results of this model are:

For a Few Agents, a Small Number of Rules, and Small Dividend Changes

- The price converges toward an equilibrium price which is close to the fundamental value.
- Trading volume is low.
- There are no bubbles, crashes, or anomalies.
- Agents follow homogeneous simple fundamentalist rules.

For a Large Number of Agents and a Large Number of Rules

- There is no convergence to an equilibrium price, and the dynamics are complex.
- The price displays occasional large deviations from the fundamental value (bubbles and crashes).
- Some of these deviations are triggered by the emergence of collectively self-fulfilling agent price-prediction rules.
- The agents become heterogeneous (adopt very different rules).
- Trading volumes fluctuate (large volumes correspond to bubbles and crashes).
- The rules evolve over time to more and more complex patterns, organized in hierarchies (rules, exceptions to rules, exceptions to exceptions, and so on).
- The successful rules are time dependent: a rule which is successful at a given time may perform poorly if reintroduced after many cycles of market coevolution.

The Lux and Lux and Marchesi Model

Lux (1998) and Lux and Marchesi (Levy et al. 1996) propose a model to endogenously explain

the heavy-tailed distribution of returns and the clustering of volatility. Both of these phenomena emerge in the Lux model as soon as one assumes that in addition to the fundamentalists, there are also chartists in the model. Lux and Marchesi (Levy et al. 1996) further divide the chartists into optimists (buyers) and pessimists (sellers). The market fluctuations are driven and amplified by the fluctuations in the various populations: chartists converting into fundamentalists, pessimists into optimists, etc.

In the Lux and Marchesi model, the stock’s fundamental value is exogenously determined. The fluctuations of the fundamental value are inputted exogenously as a white noise process in the logarithm of the value. The market price is determined by investors’ demands and by the market clearance condition.

Lux and Marchesi consider three types of traders:

- **Fundamentalists** observe the fundamental value of the stock. They anticipate that the price will eventually converge to the fundamental value, and their demand for shares is proportional to the difference between the market price and the fundamental value.
- **Chartists** look more at the present trends in the market price rather than at fundamental economic values; the chartists are divided into:
 - **Optimists** (they buy a fixed amount of shares per unit time)
 - **Pessimists** (they sell shares)

Transitions between these three groups (optimists, pessimists, fundamentalists) happen with probabilities depending on the market dynamics and on the present numbers of traders in each of the three classes:

- **The transition probabilities of chartists** depend on the majority opinion (through an “opinion index” measuring the relative number of optimists minus the relative number of pessimists) and on the actual price trend (the current time derivative of the current market price), which determines the relative profit of the various strategies.
- **The fundamentalists decide to turn into chartists** if the profits of the latter become

significantly larger than their own and vice versa (the detailed formulae used by Lux and Marchesi are inspired from the exponential transition probabilities governing statistical mechanics physical systems).

The main results of the model are:

- No long-term deviations between the current market price and the fundamental price are observed.
- The deviations from the fundamental price, which do occur, are unsystematic.
- In spite of the fact that the variations of the fundamental price are normally distributed, the variations of the market price (the market returns) are not. In particular the returns exhibit a frequency of extreme events which is higher than expected for a normal distribution. The authors emphasize the amplification role of the market that transforms the input normal distribution of the fundamental value variations into a leptokurtotic (heavy-tailed) distribution of price variation, which is encountered in the actual financial data.
- Clustering of volatility.

The authors explain the volatility clustering (and as a consequence, the leptokurticity) by the following mechanism. In periods of high volatility, the fundamental information is not very useful to insure profits, and a large fraction of the agents become chartists. The opposite is true in quiet periods when the actual price is very close to the fundamental value. The two regimes are separated by a threshold in the number of chartist agents. Once this threshold is approached (from below), large fluctuations take place which further increase the number of chartists. This destabilization is eventually dampened by the energetic intervention of the fundamentalists when the price deviates too much from the fundamental value. The authors compare this temporal instability with the on-off intermittence encountered in certain physical systems. According to Egenter et al. (1999), the fraction of chartists in the Lux Marchesi model goes to zero as the total number of traders goes to infinity, when the rest of the parameters are kept constant.

Illustration with the LLS Model

The purpose of this section is to give a more detailed “hands-on” example of the agent-based approach and to discuss some of the practical dilemmas arising when implementing this approach, by focusing on one specific model. We will focus on the so-called LLS model of the stock market (for more details and various versions of the model, see Hellthaler (1995), Kohl (1997), Levy and Levy (1996), and Levy et al. (1994, 1996, 2000)). This section is based on the presentation of the LLS model in Chap. 7 of Levy et al. (2000).

Background

Real-life investors differ in their investment behavior from the investment behavior of the idealized representative rational investor assumed in most economic and financial models. Investors differ one from the other in their preferences, their investment horizon, the information at their disposal, and their interpretation of this information. No financial economist seriously doubts these observations. However, modeling the empirically and experimentally documented investor behavior and the heterogeneity of investors is very difficult and in most cases practically impossible to do within an analytic framework. For instance, the empirical and experimental evidence suggests that most investors are characterized by Constant Relative Risk Aversion (CRRA), which implies a power (myopic) utility function (see Eq. 2 below). However, for a general distribution of returns, it is impossible to obtain an analytic solution for the portfolio optimization problem of investors with these preferences. Extrapolation of future returns from past returns, biased probability weighting, and partial deviations from rationality are also all experimentally documented but difficult to incorporate in an analytical setting. One is then usually forced to make the assumptions of rationality and homogeneity (at least in some dimension) and to make unrealistic assumptions regarding investors’ preferences, in order to obtain a model with a tractable solution. The hope in these circumstances is that the model will capture the essence of the system under investigation

and will serve as a useful benchmark, even though some of the underlying assumptions are admittedly false.

Most homogeneous rational agent models lead to the following predictions: no trading volume, zero autocorrelation of returns, and price volatility which is equal to or lower than the volatility of the “fundamental value” of the stock (defined as the present value of all future dividends; see Shiller (1981)). However, the empirical evidence is very different:

- Trading volume can be extremely heavy (Admati and Pfleiderer 1988; Karpoff 1987).
- Stock returns exhibit short-run momentum (positive autocorrelation) and long-run mean reversion (negative autocorrelation) (Fama and French 1988; Jegadeesh and Titman 1993; Levy and Lim 1998; Poterba and Summers 1988).
- Stock returns are excessively volatile relative to the dividends (Shiller 1981).

As most standard rational-representative-agent models cannot explain these empirical findings, these phenomena are known as “anomalies” or “puzzles.” Can these “anomalies” be due to elements of investors’ behavior which are unmodeled in the standard rational-representative-agent models, such as the experimentally documented deviations of investors’ behavior from rationality and/or the heterogeneity of investors? The agent-based simulation approach offers us a tool to investigate this question. The strength of the agent-based simulation approach is that since it is not restricted to the scope of analytical methods, one is able to investigate virtually any imaginable investor behavior and market structure. Thus, one can study models which incorporate the experimental findings regarding the behavior of investors and evaluate the effects of various behavioral elements on market dynamics and asset pricing.

The LLS model incorporates some of the main empirical findings regarding investor behavior, and we employ this model in order to study the effect of each element of investor behavior on asset pricing and market dynamics. We start out

with a benchmark model in which all of the investors are rational, informed, and identical, and then, one by one, we add elements of heterogeneity and deviations from rationality to the model in order to study their effects on the market dynamics.

In the benchmark model all investors are rational, informed, and identical (RII investors). This is, in effect, a “representative agent” model. The RII investors are informed about the dividend process, and they rationally act to maximize their expected utility. The RII investors make investment decisions based on the present value of future cash flows. They are essentially fundamentalists who evaluate the stock’s fundamental value and try to find bargains in the market. The benchmark model in which all investors are RII yields results which are typical of most rational-representative-agent models: in this model prices follow a random walk, there is no excess volatility of the prices relative to the volatility of the dividend process, and since all agents are identical, there is no trading volume.

After describing the properties of the benchmark model, we investigate the effects of introducing various elements of investor behavior which are found in laboratory experiments but are absent in most standard models. We do so by adding to the model a minority of investors who do not operate like the RII investors. These investors are efficient market believers (EMBs from now on). The EMBs are investors who believe that the price of the stock reflects all of the currently available information about the stock. As a consequence, they do not try to time the market or to buy bargain stocks. Rather, their investment decision is reduced to the optimal diversification problem. For this portfolio optimization, the ex-ante return distribution is required. However, since the ex-ante distribution is unknown, the EMB investors use the ex-post distribution in order to estimate the ex-ante distribution. It has been documented that in fact, many investors form their expectations regarding the future return distribution based on the distribution of past returns.

There are various ways to incorporate the investment decisions of the EMBs. This stems from the fact that there are different ways to

estimate the ex-ante distribution from the ex-post distribution. How far back should one look at the historical returns? Should more emphasis be given to more recent returns? Should some “outlier” observations be filtered out? etc. Of course, there are no clear answers to these questions, and different investors may have different ways of forming their estimation of the ex-ante return distribution (even though they are looking at the same series of historical returns). Moreover, some investors may use the objective ex-post probabilities when constructing their estimation of the ex-ante distribution, whereas others may use biased subjective probability weights. In order to build the analysis step-by-step, we start by analyzing the case in which the EMB population is homogeneous, and then introduce various forms of heterogeneity into this population.

An important issue in market modeling is that of the degree of investors’ rationality. Most models in economics and finance assume that people are fully rational. This assumption usually manifests itself as the maximization of an expected utility function by the individual. However, numerous experimental studies have shown that people deviate from rational decision-making (Thaler 1993, 1994; Tversky and Kahneman 1981, 1986, 1992). Some studies model deviations from the behavior of the rational agent by introducing a subgroup of liquidity or “noise” traders. These are traders that buy and sell stocks for reasons that are not directly related to the future payoffs of the financial asset – their motivation to trade arises from outside of the market (e.g., a “noise trader’s” daughter unexpectedly announces his/her plans to marry, and the trader sells stocks because of this unexpected need for cash). The exogenous reasons for trading are assumed random and thus lead to random or “noise” trading (see Grossman and Stiglitz 1980). The LLS model takes a different approach to the modeling of noise trading. Rather than dividing investors into the extreme categories of “fully rational” and “noise traders,” the LLS model assumes that most investors try to act as rationally as they can, but are influenced by a multitude of factors causing them to deviate to some extent from the behavior that would have been optimal from their point of view. Namely, all investors are characterized by a utility function and act to maximize their expected utility; however, some

investors may deviate to some extent from the optimal choice which maximizes their expected utility. These deviations from the optimal choice may be due to irrationality, inefficiency, liquidity constraints, or a combination of all of the above.

In the framework of the LLS model, we examine the effects of the EMBs’ deviations from rationality and their heterogeneity, relative to the benchmark model in which investors are informed, rational, and homogeneous. We find that the behavioral elements which are empirically documented, namely, extrapolation from past returns, deviation from rationality, and heterogeneity among investors, lead to all of the following empirically documented “puzzles”:

- Excess volatility
- Short-term momentum
- Longer-term return mean reversion
- Heavy trading volume
- Positive correlation between volume and contemporaneous absolute returns
- Positive correlation between volume and lagged absolute returns

The fact that all these anomalies or “puzzles,” which are hard to explain with standard rational-representative-agent models, are generated naturally by a simple model which incorporates the experimental findings regarding investor behavior and the heterogeneity of investors leads one to suspect that these behavioral elements and the diversity of investors are a crucial part of the workings of the market, and as such they cannot be “assumed away.” As the experimentally documented bounded-rational behavior and heterogeneity are in many cases impossible to analyze analytically, agent-based simulation presents a very promising tool for investigating market models incorporating these elements.

The LLS Model

The stock market consists of two investment alternatives: a stock (or index of stocks) and a bond. The bond is assumed to be a riskless asset, and the stock is a risky asset. The stock serves as a proxy for the market portfolio (e.g., the Standard & Poor’s 500 index). The extension from one risky asset to many risky assets is possible; however,

one stock (the index) is sufficient for our present analysis because we restrict ourselves to global market phenomena and do not wish to deal with asset allocation across several risky assets. Investors are allowed to revise their portfolio at given time points, i.e., we discuss a discrete time model.

The bond is assumed to be a riskless investment yielding a constant return at the end of each time period. The bond is in infinite supply and investors can buy from it as much as they wish at a given rate of r_f . The stock is in finite supply. There are N outstanding shares of the stock. The return on the stock is composed of two elements:

- (a) Capital gain: If an investor holds a stock, any rise (fall) in the price of the stock contributes to an increase (decrease) in the investor's wealth.
- (b) Dividends: The company earns income and distributes dividends at the end of each time period. We denote the dividend per share paid at time t by D_t . We assume that the dividend is a stochastic variable following a multiplicative random walk, i.e., $\tilde{D}_t = D_{t-1}(1 + \tilde{z})$, where $\tilde{z}$ is a random variable with some probability density function $f(z)$ in the range $[z_1, z_2]$ (in order to allow for a dividend cut as well as a dividend increase, we typically choose: $z_1 < 0, z_2 > 0$).

The total return on the stock in period t , which we denote by R_t is given by

$$\tilde{R}_t = \frac{\tilde{P}_t + \tilde{D}_t}{P_{t-1}}, \quad (1)$$

where $\tilde{P}_t$ is the stock price at time t .

All investors in the model are characterized by a von Neumann-Morgenstern utility function. We assume that all investors have a power utility function of the form:

$$U(W) = \frac{W^{1-\alpha}}{1-\alpha}, \quad (2)$$

where α is the risk-aversion parameter. This form of utility function implies Constant Relative Risk Aversion (CRRA). We employ the power utility function (Eq. 2) because the empirical evidence suggests that relative risk aversion is approximately constant (e.g., see Friend and Blume (1975), Gordon et al. (1972), Kroll et al. (1988), and Levy

(1994), and the power utility function is the unique utility function which satisfies the CRRA condition. Another implication of CRRA is that the optimal investment choice is independent of the investment horizon (Samuelson 1989, 1994). In other words, regardless of investors' actual investment horizon, they choose their optimal portfolio as though they are investing for a single period. The myopia property of the power utility function simplifies our analysis, as it allows us to assume that investors maximize their one-period-ahead expected utility.

We model two different types of investors: rational, informed, and identical (RII) investors and efficient market believers (EMB). These two investor types are described below.

Rational Informed Identical (RII) Investors

RII investors evaluate the "fundamental value" of the stock as the discounted stream of all future dividends and thus can also be thought of as "fundamentalists." They believe that the stock price may deviate from the fundamental value in the short run, but if it does, it will eventually converge to the fundamental value. The RII investors act according to the assumption of asymptotic convergence: if the stock price is low relative to the fundamental value, they buy in anticipation that the underpricing will be corrected and vice versa. We make the simplifying assumption that the RII investors believe that the convergence of the price to the fundamental value will occur in the next period; however, our results hold for the more general case where the convergence is assumed to occur some T periods ahead, with $T > 1$.

In order to estimate next period's return distribution, the RII investors need to estimate the distribution of next period's price, $\tilde{P}_{t+1}$, and of next period's dividend, $\tilde{D}_{t+1}$. Since they know the dividend process, the RII investors know that $\tilde{D}_{t+1} = D_t(1 + \tilde{z})$ where $\tilde{z}$ is distributed according to $f(z)$ in the range $[z_1, z_2]$. The RII investors employ Gordon's dividend stream model in order to calculate the fundamental value of the stock:

$$p_{t+1}^f = \frac{E_{t+1}[\tilde{D}_{t+2}]}{k - g}, \quad (3)$$

where the superscript f stands for the *fundamental* value, $E_{t+1}[\tilde{D}_{t+2}]$ is the dividend corresponding

to time $t + 2$ as expected at time $t + 1$, k is the discount factor or the expected rate of return demanded by the market for the stock, and g is the *expected* growth rate of the dividend, i.e.,

$$g = E(\tilde{z}) = \int_{z_1}^{z_2} f(z)zdz.$$

The RII investors believe that the stock price may temporarily deviate from the fundamental value; however, they also believe that the price will eventually converge to the fundamental value. For simplification we assume that the RII investors believe that the convergence to the fundamental value will take place next period. Thus, the RII investors estimate $P_t + 1$ as

$$P_{t+1} = P_{t+1}^f.$$

The expectation at time $t + 1$ of $\tilde{D}_{t+2}$ depends on the realized dividend observed at $t + 1$:

$$E_{t+1}[\tilde{D}_{t+2}] = D_{t+1}(1 + g).$$

Thus, the RII investors believe that the price at $t + 1$ will be given by

$$P_{t+1} = P_{t+1}^f = \frac{D_{t+1}(1 + g)}{k - g}.$$

At time t , D_t is known, but D_{t+1} is not; therefore P_{t+1}^f is also not known with certainty at time t .

However, given D_t , the RII investors know the distribution of $\tilde{D}_{t+1}$:

where $\tilde{z}$ is distributed according to the known $f(z)$. The realization of $\tilde{D}_{t+1}$ determines P_{t+1}^f . Thus, at time t , RII investors believe that P_{t+1} is a random variable given by

$$\tilde{P}_{t+1} = \tilde{P}_{t+1}^f = \frac{D_t(1 + \tilde{z})(1 + g)}{k - g}.$$

Notice that the RII investors face uncertainty regarding next period's price. In our model we assume that the RII investors are certain about the dividend growth rate g , the discount factor k , and the fact that the price will converge to the

fundamental value next period. In this framework the only source of uncertainty regarding next period's price stems from the uncertainty regarding next period's dividend realization. More generally, the RII investors' uncertainty can result from uncertainty regarding any one of the above factors or a combination of several of these factors. Any mix of these uncertainties is possible to investigate in the agent-based simulation framework, but very hard, if not impossible, to incorporate in an analytic framework. As a consequence of the uncertainty regarding next period's price and of their risk aversion, the RII investors do not buy an infinite number of shares even if they perceive the stock as underpriced. Rather, they estimate the stock's next period's return distribution and find the optimal mix of the stock and the bond which maximizes their expected utility. The RII investors estimate next period's return on the stock as

$$\begin{aligned} \tilde{R}_{t+1} &= \frac{\tilde{P}_{t+1} + \tilde{D}_{t+1}}{P_t} \\ &= \frac{\frac{D_t(1 + \tilde{z})(1 + g)}{k - g} + D_t(1 + \tilde{z})}{P_t}, \end{aligned} \quad (4)$$

where $\tilde{z}$, the next year's growth in the dividend, is the source of uncertainty. The demands of the RII investors for the stock depend on the price of the stock. For any *hypothetical* price P_h , investors calculate the proportion of their wealth x they should invest in the stock in order to maximize their expected utility. The RII investor i believes that if he/she invests a proportion x of his/her wealth in the stock at time t , then at time $t + 1$ his/her wealth will be

$$\tilde{W}_{t+1}^i = W_h^i \left[(1 - x)(1 + r_f) + x\tilde{R}_{t+1} \right], \quad (5)$$

where $\tilde{R}_{t+1}$ is the return on the stock, as given by Eq. 1, and is the wealth of investor W_h^i at time t given that the stock price at time t is P_h .

If the price in period t is the hypothetical price P_h , the $t + 1$ expected utility of investor i is the following function of his/her investment proportion in the stock, x :

$$EU(\tilde{W}_{t+1}^i) = EU\left(W_h^i \left[(1-x)(1+r_f) + x\tilde{R}_{t+1} \right]\right). \quad (6)$$

Substituting $\tilde{R}_{t+1}$ from Eq. 4, using the power utility function (Eq. 2), and substituting the hypothetical price P_h for P_t , the expected utility becomes the following function of x :

$$EU(\tilde{W}_{t+1}^i) = \frac{(W_h^i)^{1-\alpha}}{1-\alpha} \int_{z_1}^{z_2} \left[(1-x)(1+r_f) + x \left(\frac{D_t(1+z)(1+g)}{k-g} + D_t(1+z) \right) \frac{1}{P_h} \right]^{1-\alpha} f(z) dz, \quad (7)$$

where the integration is over all possible values of z . In the agent-based simulation framework, this expression for the expected utility, and the optimal investment proportion x , can be solved numerically for any general choice of distribution $f(z)$. For the sake of simplicity, we restrict the present analysis to the case where $\tilde{z}$ is distributed uniformly in the range $[z_1, z_2]$. This simplification leads to the following expression for the expected utility:

$$EU(\tilde{W}_{t+1}^i) = \frac{(W_h^i)^{1-\alpha}}{(1-\alpha)(2-\alpha)} \frac{1}{(z_2 - z_1)} \left\{ \left[(1-x)(1+r_f) + \frac{x}{P_h} \left(\frac{k+1}{k-g} \right) D_t(1+z_2) \right]^{(2-\alpha)} - \left[(1-x)(1+r_f) + \frac{x}{P_h} \left(\frac{k+1}{k-g} \right) D_t(1+z_1) \right]^{(2-\alpha)} \right\}. \quad (8)$$

For any hypothetical price P_h , each investor (numerically) finds the optimal proportion x_h which maximizes his/her expected utility given by Eq. 8. Notice that the optimal proportion, x_h , is independent of the wealth, W_h^i . Thus, if all RII investors have the same degree of risk aversion, α , they will have the same optimal investment proportion in the stock, regardless of their wealth. The number of shares demanded by investor i at the hypothetical price P_h is given by

$$N_h^i(P_h) = \frac{x_h^i(P_h) W_h^i(P_h)}{P_h}. \quad (9)$$

Efficient Market Believers (EMB)

The second type of investors in the LLS model is EMBs. The EMBs believe in market efficiency – they believe that the stock price accurately reflects the stock’s fundamental value. Thus, they do not try to time the market or to look for “bargain” stocks. Rather, their investment decision is reduced to the optimal diversification between the stock and the bond. This diversification decision requires the ex-ante return distribution for the stock, but as the ex-ante distribution is not available, the EMBs assume that the process generating the returns is fairly stable, and they employ the ex-post distribution of stock returns in order to estimate the ex-ante return distribution.

Different EMB investors may disagree on the optimal number of ex-post return observations that should be employed in order to estimate the ex-ante return distribution. There is a trade-off between using more observations for better statistical inference and using a smaller number of only more recent observations, which are probably more representative of the ex-ante distribution. As in reality, there is no “recipe” for the optimal number of observations to use. EMB investor i believes that the m^i most recent returns on the stock are the best estimate of the ex-ante distribution. Investors create an estimation of the ex-ante return distribution by assigning an equal probability to each of the m^i most recent return observations:

$$\text{Prob}^i(\tilde{R}_{t+1} = R_{t-j}) = \frac{1}{m^i} \text{ for } j = 1, \dots, m^i. \quad (10)$$

The expected utility of EMB investor i is given by

$$EU(W_{t+1}^i) = \frac{(W_h^i)^{1-\alpha}}{1-\alpha} \frac{1}{m^i} \sum_{j=1}^{m^i} \left[(1-x)(1+r_f) + x R_{t-j} \right]^{1-\alpha}, \quad (11)$$

where the summation is over the set of m^i most recent ex-post returns, x is the proportion of wealth invested in the stock, and as before W_h^i is the wealth of investor i at time t given that the stock price at time t is P_h . Notice that W_h^i does not

change the optimal diversification policy, i.e., x . Given a set of m^i past returns, the optimal portfolio for the EMB investor i is an investment of a proportion x^{*i} in the stock and $(1-x^{*i})$ in the bond, where x^{*i} is the proportion which maximizes the above expected utility (Eq. 11) for investor i . Notice that x^{*i} generally cannot be solved for analytically. However, in the agent-based simulation framework, this does not constitute a problem, as one can find x^{*i} numerically.

Deviations from Rationality

Investors who are efficient market believers, and are rational, choose the investment proportion x^* which maximizes their expected utility. However, many empirical studies have shown that the behavior of investors is driven not only by rational expected utility maximization but by a multitude of other factors (e.g., see Samuelson (1994), Thaler (1993, 1994), and Tversky and Kahneman (1981, 1986)). Deviations from the optimal rational investment proportion can be due to the cost of resources which are required for the portfolio optimization – time, access to information, computational power, etc. – or due to exogenous events (e.g., an investor plans to revise his/her portfolio, but gets distracted because his/her car breaks down). We assume that the different factors causing the investor to deviate from the optimal investment proportion x^* are random and uncorrelated with each other. By the central limit theorem, the aggregate effect of a large number of random uncorrelated influences is a normally distributed random influence, or “noise.” Hence, we model the effect of all the factors causing the investor to deviate from his/her optimal portfolio by adding a normally distributed random variable to the optimal investment proportion. To be more specific, we assume

$$x^i = x^{*i} + \tilde{\varepsilon}^i, \quad (12)$$

where $\tilde{\varepsilon}^i$ is a random variable drawn from a truncated normal distribution with mean zero and standard deviation σ . Notice that noise is investor specific; thus, $\tilde{\varepsilon}^i$ is drawn separately and independently for each investor.

The noise can be added to the decision-making of the RII investors, the EMB investors, or to both.

The results are not much different with these various approaches. Since the RII investors are taken as the benchmark of rationality, in this entry we add the noise only to the decision-making of the EMB investors.

Market Clearance

The number of shares demanded by each investor is a monotonically decreasing function of the hypothetical price P_h (see Levy et al. 2000). As the total number of outstanding shares is N , the price of the stock at time t is given by the market clearance condition: P_t is the unique price at which the total demand for shares is equal to the total supply, N :

$$\sum_i N_h^i(P_t) = \sum_i \frac{x_h(P_t) W_h^i(P_t)}{P_t} = N, \quad (13)$$

where the summation is over all the investors in the market, RII investors as well as EMB investors.

Agent-Based Simulation

The market dynamics begin with a set of initial conditions which consist of an initial stock price P_0 , an initial dividend D_0 , the wealth and number of shares held by each investor at time $t = 0$, and an initial “history” of stock returns. As will become evident, the general results do not depend on the initial conditions. At the first period ($t = 1$), interest is paid on the bond, and the time 1 dividend $\tilde{D}_1 = D_0(1 + \tilde{z})$ is realized and paid out. Then investors submit their demand orders, $N_h^i(P_h)$, and the market clearing price P_1 is determined. After the clearing price is set, the new wealth and number of shares held by each investor are calculated. This completes one time period. This process is repeated over and over, as the market dynamics develop.

We would like to stress that even the simplified benchmark model, with only RII investors, is impossible to solve analytically. The reason for this is that the optimal investment proportion, $x_h(P_h)$, cannot be calculated analytically. This problem is very general and it is encountered with almost any choice of utility function and

distribution of returns. One important exception is the case of a negative exponential utility function and normally distributed returns. Indeed, many models make these two assumptions for the sake of tractability. The problem with the assumption of negative exponential utility is that it implies Constant Absolute Risk Aversion (CARA), which is very unrealistic, as it implies that investors choose to invest the same dollar amount in a risky prospect *independent of their wealth*. This is not only in sharp contradiction to the empirical evidence but also excludes the investigation of the two-way interaction between wealth and price dynamics, which is crucial to the understanding of the market.

Thus, one contribution of the agent-based simulation approach is that it allows investigation of models with realistic assumptions regarding investors' preferences. However, the main contribution of this method is that it permits us to investigate models which are much more complex (and realistic) than the benchmark model, in which all investors are RII. With the agent-based simulation approach, one can study models incorporating the empirically and experimentally documented investors' behavior and the heterogeneity of investors.

Results of the LLS Model

We begin by describing the benchmark case where all investors are rational and identical. Then we introduce to the market EMB investors and investigate their affects on the market dynamics.

Benchmark Case: Fully Rational and Identical Agents

In this benchmark model all investors are RII: rational, informed, and identical. Thus, it is not surprising that the benchmark model generates market dynamics which are typical of homogeneous rational agent models.

No Volume All investors in the model are identical; they therefore always agree on the optimal proportion to invest in the stock. As a consequence, all the investors always achieve the same return on their portfolio. This means that at

any time period, the ratio between the wealth of any two investors is equal to the ratio of their initial wealths, i.e.,:

$$\frac{W_t^i}{W_t^j} = \frac{W_0^i}{W_0^j}. \quad (14)$$

As the wealth of investors is always in the same proportion, and as they always invest the same fraction of their wealth in the stock, the number of shares held by different investors is also always in the same proportion:

$$\frac{N_t^i}{N_t^j} = \frac{\frac{x_i W_t^i}{P_t}}{\frac{x_j W_t^j}{P_t}} = \frac{W_t^i}{W_t^j} = \frac{W_0^i}{W_0^j}. \quad (15)$$

Since the total supply of shares is constant, this implies that each investor always holds the same number of shares and there is *no trading volume* (the number of shares held may vary from one investor to another as a consequence of different initial endowments).

Log Prices Follow a Random Walk In the benchmark model all investors believe that next period's price will converge to the fundamental value given by the discounted dividend model (Eq. 3). Therefore, the actual stock price is always close to the fundamental value. The fluctuations in the stock price are driven by fluctuations in the fundamental value, which in turn are driven by the fluctuating dividend realizations. As the dividend fluctuations are (by assumption) uncorrelated over time, one would expect that the price fluctuations will also be uncorrelated. To verify this intuitive result, we examine the return autocorrelations in simulations of the benchmark model.

Let us turn to the simulation of the model. We first describe the parameters and initial conditions used in the simulation and then report the results. We simulate the benchmark model with the following parameters:

- Number of investors = 1,000.
- Risk-aversion parameter $\alpha = 1.5$. This value roughly conforms with the estimate of the risk-

aversion parameter found empirically and experimentally.

- Number of shares = 10,000.
- We take the time period to be a quarter, and accordingly we choose:
- Riskless interest rate $r_f = 0.01$.
- Required rate of return on stock $k = 0.04$.
- Maximal one-period dividend decrease $z_1 = -0.07$.
- Maximal one-period dividend growth $z_2 = 0.10$.
- $\tilde{z}$ is uniformly distributed between these values. Thus, the average dividend growth rate is $g = (z_1 + z_2)/2 = 0.015$.

Initial conditions: Each investor is endowed at time $t = 0$ with a total wealth of \$1,000, which is composed of 10 shares worth an initial price of \$50 per share and \$500 in cash. The initial quarterly dividend is set at \$0.5 (for an annual dividend yield of about 4%). As will soon become evident, the dynamics are not sensitive to the particular choice of initial conditions.

Figure 1 shows the price dynamics in a typical simulation with these parameters (simulations with the same parameters differ one from the other because of the different random dividend realizations). Notice that the vertical axis in this figure is logarithmic. Thus, the roughly constant slope implies an approximately exponential price growth or an approximately constant *average* return.

The prices in this simulation seem to fluctuate randomly around the trend. However, Fig. 1 shows only one simulation. In order to have a more rigorous analysis, we perform many independent simulations and employ statistical tools. Namely, for each simulation we calculate the autocorrelation of returns. We perform a univariate regression of the return in time t on the return on time $t - j$:

$$R_t = \alpha_j + \beta_j R_{t-j} + \varepsilon,$$


where R_t is the return in period t and j is the lag. The autocorrelation of returns for lag j is defined as

$$\rho_j = \frac{\text{cov}(R_t, R_{t-j})}{\hat{\sigma}^2(R)},$$


and it is estimated by $\hat{\beta}$. We calculate the autocorrelation for different lags, $j = 1, \dots, 40$. Figure 2 shows the average autocorrelation as a function of the lag, calculated over 100 independent simulations. It is evident both from the figure that the returns are uncorrelated in the benchmark model, conforming with the random-walk hypothesis.

No Excess Volatility Since the RII investors believe that the stock price will converge to the fundamental value next period, in the benchmark model, prices are always close to the fundamental value given by the discounted dividend stream. Thus, we do not expect prices to be more volatile

**Agent-Based
Computational
Economics, Fig. 1** Price
dynamics in the benchmark
model



**Agent-Based
Computational
Economics, Fig. 2** Return
autocorrelation in
benchmark model



than the value of the discounted dividend stream. For a formal test of excess volatility, we follow the technique in Shiller (1981). For each time period we calculate the actual price, P_t , and the fundamental value of discounted dividend stream, p_t^f , as in Eq. 3. Since prices follow an upward trend, in order to have a meaningful measure of the volatility, we must detrend these price series. Following Shiller, we run the regression:

$$\ln P_t = bt + c + \varepsilon_t, \quad (16)$$

in order to find the average exponential price growth rate (where b and c are constants). Then, we define the detrended price as: $p_t = P_t / e^{\widehat{bt}}$. Similarly, we define the detrended value of the discounted dividend stream p_t^f and compare $\sigma(p_t)$ with $\sigma(p_t^f)$. For 100- to 1,000-period simulations, we find an average $\sigma(p_t)$ of 22.4 and an average $\sigma(p_t^f)$ of 22.9. As expected, the actual price and the fundamental value have almost the same volatility.

To summarize the results obtained for the benchmark model, we find that when all investors are assumed to be rational, informed, and identical, we obtain results which are typical of rational-representative-agent models: no volume, no return autocorrelations, and no excess volatility. We next turn to examine the effect of introducing into the market EMB investors, which model empirically and experimentally documented elements of investors' behavior.

The Introduction of a Small Minority of EMB Investors

In this section we will show that the introduction of a small minority of heterogeneous EMB investors generates many of the empirically observed market "anomalies" which are absent in the benchmark model and indeed, in most other rational-representative-agent models. We take this as strong evidence that the "nonrational" elements of investor behavior which are documented in experimental studies and the heterogeneity of investors, both of which are incorporated in the LLS model, are crucial to understanding the dynamics of the market.

In presenting the results of the LLS model with EMB investors, we take an incremental approach. We begin by describing the results of a model with a small subpopulation of *homogeneous* EMB believers. This model produces the abovementioned market "anomalies"; however, it produces unrealistic cyclic market dynamics. Thus, this model is presented both for analyzing the source of the "anomalies" in a simplified setting and as a reference point with which to compare the dynamics of the model with a *heterogeneous* EMB believer population.

We investigate the effects of investors' heterogeneity by first analyzing the case in which there are two types of EMBs. The two types differ in the method they use to estimate the ex-ante return distribution. Namely, the first type looks at the set of the last m_1 ex-post returns, whereas the second type looks at the set of the last m_2 ex-post returns. It turns out that the dynamics in

this case are much more complicated than a simple “average” between the case where all EMB investors have m_1 and the case where all EMB investors have m_2 . Rather, there is a complex nonlinear interaction between the two EMB subpopulations. This implies that the heterogeneity of investors is a very important element determining the market dynamics, an element which is completely absent in representative-agent models.

Finally, we present the case where there is an entire spectrum of EMB investors differing in the number of ex-post observations they take into account when estimating the ex-ante distribution. This general case generates very realistic-looking market dynamics with all of the abovementioned market anomalies.


Homogeneous Subpopulation of EMBs

When a very small subpopulation of EMB investors is introduced to the benchmark LLS model, the market dynamics change dramatically. Figure 3 depicts a typical price path in a simulation of a market with 95% RII investors and 5% EMB investors. The EMB investors have $m = 10$ (i.e., they estimate the ex-ante return distribution by observing the set of the last 10 ex-post returns). σ , the standard deviation of the random noise affecting the EMBs’ decision-making, is taken as 0.2. All investors, RII and EMB alike, have the same risk-aversion parameter $\alpha = 1.5$ (as before). In the first 150 trading periods, the price dynamics

look very similar to the typical dynamics of the benchmark model. However, after the first 150 or so periods, the price dynamics change. From this point onwards the market is characterized by periodic booms and crashes. Of course, Fig. 3 describes only one simulation. However, as will become evident shortly, different simulations with the same parameters may differ in detail, but the pattern is general: at some stage (not necessarily after 150 periods), the EMB investors induce cyclic price behavior. It is quite astonishing that such a small minority of only 5% of the investors can have such a dramatic impact on the market.

In order to understand the periodic booms and crashes, let us focus on the behavior of the EMB investors. After every trade, the EMB investors revise their estimation of the ex-ante return distribution, because the set of ex-post returns they employ to estimate the ex-ante distribution changes. Namely, investors add the latest return generated by the stock to this set and delete the oldest return from this set. As a result of this update in the estimation of the ex-ante distribution, the optimal investment proportion x^* changes, and EMB investors revise their portfolios at next period’s trade. During the first 150 or so periods, the informed investors control the dynamics and the returns fluctuate randomly (as in the benchmark model). As a consequence, the investment proportion of the EMB investors also fluctuates irregularly. Thus, during the first 150 periods, the EMB

**Agent-Based
Computational
Economics, Fig. 3** Five percent of investors are efficient market believers, 95% rational informed investors



investors do not effect the dynamics much. However, at point **a**, the dynamics change qualitatively (see Fig. 3). At this point, a relatively high dividend is realized, and as a consequence, a relatively high return is generated. This high return leads the EMB investors to increase their investment proportion in the stock at the next trading period. This increased demand of the EMB investors is large enough to effect next period's price, and thus a second high return is generated. Now the EMB investors look at a set of ex-post returns with two high returns, and they increase their investment proportion even further. Thus, a positive feedback loop is created.

Notice that as the price goes up, the informed investors realize that the stock is overvalued relative to the fundamental value P^f and they decrease their holdings in the stock. However, this effect does not stop the price increase and break the feedback loop because the EMB investors continue to buy shares aggressively. The positive feedback loop pushes the stock price further and further up to point **b**, at which the EMBs are invested 100 % in the stock. At point **b**, the positive feedback loop "runs out of gas." However, the stock price remains at the high level because the EMB investors remain fully invested in the stock (the set of past $m = 10$ returns includes at this stage the very high returns generated during the "boom" – segment **a–b** in Fig. 3).


When the price is at the high level (segment **b–c**), the dividend yield is low, and as a consequence, the returns are generally low. As time goes by and we move from point **b** toward point **c**, the set of $m = 10$ last returns gets filled with low returns. Despite this fact, the extremely high returns generated in the boom are also still in this set, and they are high enough to keep the EMB investors fully invested. However, 10 periods after the boom, these extremely high returns are pushed out of the set of relevant ex-post returns. When this occurs, at point **c**, the EMB investors face a set of low returns, and they cut their investment proportion in the stock sharply. This causes a dramatic crash (segment **c–d**). Once the stock price goes back down to the "fundamental" value, the informed investors come back into the picture. They buy back the stock and stop the crash.

The EMB investors stay away from the stock as long as the ex-post return set includes the terrible return of the crash. At this stage the informed investors regain control of the dynamics and the stock price remains close to its fundamental value. Ten periods after the crash, the extremely negative return of the crash is excluded from the ex-post return set, and the EMB investors start increasing their investment proportion in the stock (point **e**). This drives the stock price up, and a new boom-crash cycle is initiated. This cycle repeats itself over and over almost periodically.

Figure 3 depicts the price dynamics of a single simulation. One may therefore wonder how general the results discussed above are. Figure 4 shows two more simulations with the same parameters but different dividend realizations. It is evident from this figure that although the simulations vary in detail (because of the different dividend realizations), the overall price pattern with periodic boom-crash cycles is robust.

Although these dynamics are very unrealistic in terms of the periodicity, and therefore the predictability of the price, they do shed light on the mechanism generating many of the empirically observed market phenomena. In the next section, when we relax the assumption that the EMB population is homogeneous with respect to m , the price is no longer cyclic or predictable, yet the mechanisms generating the market phenomena are the same as in this homogeneous EMB population case. The homogeneous EMB population case generates the following market phenomena.

Heavy Trading Volume As explained above, shares change hands continuously between the RII investors and the EMB investors. When a "boom" starts, the RII investors observe higher ex-post returns and become more optimistic, while the EMB investors view the stock as becoming overpriced and become more pessimistic. Thus, at this stage the EMBs buy most of the shares from the RIIs. When the stock crashes, the opposite is true: the EMBs are very pessimistic, but the RII investors buy the stock once it falls back to the fundamental value. Thus, there is substantial trading volume in this market. The average trading volume in a typical



Agent-Based Computational Economics, Fig. 4 Two more simulations – same parameters as Fig. 3, different dividend realizations

simulation is about 1,000 shares per period, which are 10 % of the total outstanding shares.

Autocorrelation of Returns The cyclic behavior of the price yields a very definite return autocorrelation pattern. The autocorrelation pattern is depicted graphically in Fig. 5. The autocorrelation pattern is directly linked to the length of the price cycle, which in turn are determined by m . Since the moving window of ex-post returns used to estimate the ex-ante distribution is $m = 10$ periods long, the price cycles are typically a little longer than 20 periods long: a cycle consists of the positive feedback loop (segment **a, b** in Fig. 3) which is about 2–3 periods long, the upper plateau (segment **b, c** in Fig. 3) which is about 10 periods long, the crash that occurs during one or two periods, and the lower plateau (segment **d, e** in Fig. 3) which is again about 10 periods long, for a total of about 23–25 periods. Thus, we expect positive autocorrelation for lags of about 23–25 periods, because this is the lag between one point and the corresponding point in the next (or previous) cycle. We also expect negative autocorrelation for lags of about 10–12 periods, because this is the lag between a boom and the following (or previous) crash and vice versa. This is precisely the pattern we observe in Fig. 5.

Excess Volatility The EMB investors induce large deviations of the price from the fundamental value. Thus, price fluctuations are caused not only by dividend fluctuations (as the standard theory suggests) but also by the endogenous market dynamics driven by the EMB investors. This “extra” source of fluctuations causes the price to be more volatile than the fundamental value P^f .


Indeed, for 100- to 1,000-period independent simulations with 5 % EMB investors, we find an average $\sigma(p_t)$ of 46.4 and an average $\sigma(p_t^f)$ of 30.6; That is, we have excess volatility of about 50 %.

As a first step in analyzing the effects of heterogeneity of the EMB population, in the next section we examine the case of two types of EMB investors. We later analyze a model in which there is a full spectrum of EMB investors.

Two Types of EMBs

One justification for using a representative agent in economic modeling is that although investors are heterogeneous in reality, one can model their collective behavior with one representative or “average” investor. In this section we show that this is generally not true. Many aspects of the dynamics result from the nonlinear interaction between different investor types. To illustrate this point, in this section we analyze a very simple

**Agent-Based
Computational
Economics, Fig. 5** Return
autocorrelation 5 %, efficient market believers,
 $m = 10$




case in which there are only two types of EMB investors: one with $m = 5$ and the other with $m = 15$. Each of these two types consists of 2 % of the investor population, and the remaining 96 % are informed investors. The representative agent logic may tempt us to think that the resulting market dynamics would be similar to that of one “average” investor, i.e., an investor with $m = 10$. Figure 6 shows that this is clearly not the case. Rather than seeing periodic cycles of about 23–25 periods (which correspond to the average m of 10, as in Fig. 3), we see an irregular pattern. As before, the dynamics are first dictated by the informed investors. Then, at point **a**, the EMB investors with $m = 15$ induce cycles which are about 30 periods long. At point **b** there is a transition to shorter cycles induced by the $m = 5$ population, and at point **c** there is another transition back to longer cycles. What is going on?


These complex dynamics result from the non-linear interaction between the different subpopulations. The transitions from one price pattern to another can be partly understood by looking at the wealth of each subpopulation. Figure 7 shows the proportion of the total wealth held by each of the two EMB populations (the remaining proportion is held by the informed investors). As seen in Fig. 7, the cycles which start at point **a** are dictated by the $m = 15$ rather than the $m = 5$ population, because at this stage the $m = 15$ population controls more of the wealth than the $m = 5$ population. However, after 3 cycles (at point **b**), the picture is reversed. At this point

the $m = 5$ population is more powerful than the $m = 15$ population, and there is a transition to shorter boom-crash cycles. At point **c** the wealth of the two subpopulations is again almost equal, and there is another transition to longer cycles. Thus, the complex price dynamics can be partly understood from the wealth dynamics. But how are the wealth dynamics determined? Why does the $m = 5$ population become wealthier at point **b**, and why does it lose most of this advantage at point **c**? It is obvious that the wealth dynamics are influenced by the price dynamics; thus, there is a complicated two-way interaction between the two. Although this interaction is generally very complex, some principle ideas about the mutual influence between the wealth and price patterns can be formulated. For example, a population that becomes dominant and dictates the price dynamics typically starts underperforming, because it affects the price with its actions. This means pushing the price up when buying, and therefore buying high, and pushing the price down when selling. However, a more detailed analysis must consider the specific investment strategy employed by each population. For a more comprehensive analysis of the interaction between heterogeneous EMB populations, see Levy et al. (1996).

The two-EMB-population model generates the same market phenomena as did the homogeneous population case: heavy trading volume, return autocorrelations, and excess volatility. Although the price pattern is much less regular



Agent-Based Computational Economics, Fig. 6 Two percent EMB $m = 5$, 2 % EMB $m = 15$, 96 % RII



Agent-Based Computational Economics, Fig. 7 Proportion of the total wealth held by the two EMB populations

in the two-EMB-population case, there still seems to be a great deal of predictability about the prices. Moreover, the booms and crashes generated by this model are unrealistically dramatic and frequent. In the next section we analyze a model with a continuous spectrum of EMB investors. We show that this fuller heterogeneity

of investors leads to very realistic price and volume patterns.

Full Spectrum of EMB Investors

Up to this point we have analyzed markets with at most three different subpopulations (one RII population and two EMB populations). The market

dynamics we found displayed the empirically observed market anomalies, but they were unrealistic in the magnitude, frequency, and semi-predictability of booms and crashes. In reality, we would expect not only two or three investor types, but rather an entire spectrum of investors. In this section we consider a model with a full spectrum of different EMB investors. It turns out that *more is different*. When there is an entire range of investors, the price dynamics become realistic: booms and crashes are not periodic or predictable, and they are also less frequent and dramatic. At the same time, we still obtain all of the market anomalies described before.


In this model each investor has a different number of ex-post observations which he/she utilizes to estimate the ex-ante distribution. Namely, investor i looks at the set of the m^i most recent returns on the stock, and we assume that m^i is distributed in the population according to a truncated normal distribution with average $\bar{m}$ and standard deviation σ_m (as $m \leq 0$ is meaningless, the distribution is truncated at $m = 0$).

Figure 8 shows the price pattern of a typical simulation of this model. In this simulation 90 % of the investors are RII, and the remaining 10 % are heterogeneous EMB investors with $\bar{m} = 40$


and $\sigma_m = 10$. The price pattern seems very realistic with “smoother” and more irregular cycles. Crashes are dramatic, but infrequent and unpredictable.

The heterogeneous EMB population model generates the following empirically observed market phenomena:

Return Autocorrelation: Momentum and Mean Reversion In the heterogeneous EMB population model, trends are generated by the same positive feedback mechanism that generated cycles in the homogeneous case: high (low) returns tend to make the EMB investors more (less) aggressive, this generates more high (low) returns, etc. The difference between the two cases is that in the heterogeneous case, there is a very complicated interaction between all the different investor subpopulations, and as a result there are no distinct regular cycles, but rather, smoother and more irregular trends. There is no single cycle length – the dynamics are a combination of many different cycles. This makes the autocorrelation pattern also smoother and more continuous. The return autocorrelations in the heterogeneous model are shown in Fig. 9. This autocorrelation pattern conforms with the empirical findings. In



Agent-Based Computational Economics, Fig. 8 Spectrum of heterogeneous EMB investors (10 % EMB investors, 90 % RII investors)



Agent-Based Computational Economics, Fig. 9 Return autocorrelation – heterogeneous EMB population

the short run (lags 1–4), the autocorrelation is positive – this is the empirically documented phenomenon known as momentum: in the short run, high returns tend to be followed by more high returns, and low returns tend to be followed by more low returns. In the longer run (lags 5–13), the autocorrelation is negative, which is known as mean reversion. For even longer lags the autocorrelation eventually tends to be zero. The short-run momentum, longer-run mean reversion, and eventual diminishing autocorrelation create the general “U shape” which is found in empirical studies (Fama and French 1988; Jegadeesh and Titman 1993; Poterba and Summers 1988) and which is seen in Fig. 9.

Excess Volatility The price level is generally determined by the fundamental value of the stock. However, as in the homogeneous EMB population case, the EMB investors occasionally induce temporary departures of the price away from the fundamental value. These temporary departures from the fundamental value make the price more volatile than the fundamental value. Following Shiller’s methodology we define the detrended price, p , and fundamental value, p^f . Averaging over 100 independent simulations, we find $\sigma(p) = 27.1$ and $\sigma(p^f)$, which is an excess volatility of 41 %.

Heavy Volume As investors in our model have different information (the informed investors know the dividend process, while the EMB investors do not) and different ways of interpreting the information (EMB investors with different memory spans have different estimations regarding the ex-ante return distribution), there is a high level of trading volume in this model. The average trading volume in this model is about 1,700 shares per period (17 % of the total outstanding shares). As explained below, the volume is positively correlated with contemporaneous and lagged absolute returns.

Volume Is Positively Correlated with Contemporaneous and Lagged Absolute Returns

Investors revise their portfolios as a result of changes in their beliefs regarding the future return distribution. The changes in the beliefs can be due to a change in the current price, to a new dividend realization (in the case of the informed investors), or to a new observation of an ex-post return (in the case of the EMB investors). If all investors change their beliefs in the same direction (e.g., if everybody becomes more optimistic), the stock price can change substantially with almost no volume – everybody would like to increase the proportion of the stock in his/her portfolio, this will push the price up, but a very small number of shares will change hands. This scenario would lead to zero or perhaps even negative correlation between the magnitude of

the price change (or return) and the volume. However, the typical scenario in the LLS model is different. Typically, when a positive feedback trend is induced by the EMB investors, the opinions of the informed investors and the EMB investors change in opposite directions. The EMB investors see a trend of rising prices as a positive indication about the ex-ante return distribution, while the informed investors believe that the higher the price level is above the fundamental value, the more overpriced the stock is and the harder it will eventually fall. The exact opposite holds for a trend of falling prices. Thus, price trends are typically interpreted differently by the two investor types and therefore induce heavy trading volume. The more pronounced the trend, the more likely it is to lead to heavy volume and, at the same time, to large price changes which are due to the positive feedback trading on behalf of the EMB investors.

This explains not only the positive correlation between volume and contemporaneous absolute rates of return but also the positive correlation between volume and lagged absolute rates of return. The reason is that the behavior of the EMB investors induces short-term positive return autocorrelation, or momentum (see above), that is, a large absolute return this period is associated not only with high volume but also with a large absolute return next period and therefore with high volume next period. In other words, when there is a substantial price increase (decrease), EMB investors become more (less) aggressive and the opposite happens to the informed traders. As we have seen before, when a positive feedback loop is started, the EMB investors are more dominant in determining the price, and therefore another large price increase (decrease) is expected next period. This large price change is likely to be associated with heavy trading volume as the opinions of the two populations diverge. Furthermore, this large increase (decrease) is expected to make the EMB investors even more optimistic (pessimistic) leading to another large price increase (decrease) and heavy volume next period.

In order to verify this relationship quantitatively, we regress volume on contemporaneous and lagged absolute rates of return for 100 independent simulations. We run the regressions:

$$\begin{aligned} V_t &= \alpha + \beta_C |R_t - 1| + \varepsilon_t \quad \text{and} \\ V_t &= \alpha + \beta_L |R_{t-1} - 1| + \varepsilon_t, \end{aligned} \quad (17)$$

where V_t is the volume at time t and R_t is the total return on the stock at time t and the subscripts C and L stand for contemporaneous and lagged. We find an average value of 870 for $\hat{\beta}_C$ with an average t -value of 5.0 and an average value of 886 $\hat{\beta}_L$ for with an average t -value of 5.1.

Discussion of the LLS Results

The LLS model is an agent-based simulation model of the stock market which incorporates some of the fundamental experimental findings regarding the behavior of investors. The main non-standard assumption of the model is that there is a small minority of investors in the market who are uninformed about the dividend process and who believe in market efficiency. The investment decision of these investors is reduced to the optimal diversification between the stock and the bond.

The LLS model generates many of the empirically documented market phenomena which are hard to explain in the analytical rational-representative-agent framework. These phenomena are:

- Short-term momentum
- Longer-term mean reversion
- Excess volatility
- Heavy trading volume
- Positive correlation between volume and contemporaneous absolute returns
- Positive correlation between volume and lagged absolute returns
- Endogenous market crashes

The fact that so many “puzzles” are explained with a simple model built on a small number of empirically documented behavioral elements leads us to suspect that these behavioral elements are very important in understanding the workings of the market. This is especially true in light of the observations that a very small minority of the nonstandard bounded-rational investors can have a dramatic influence on the market and that these investors are not wiped out by the majority of rational investors.

Summary and Future Directions

Standard economic models typically describe a world of homogeneous rational agents. This approach is the foundation of most of our present-day knowledge in economic theory. With the agent-based simulation approach, we can investigate a much more complex and “messy” world with different agent types, who employ different strategies to try to survive and prosper in a market with structural uncertainty. Agents can learn over time, from their own experience and from their observation about the performance of other agents. They coevolve over time and as they do so, the market dynamics change continuously. This is a worldview closer to biology than it is to the “clean” realm of physical laws which classical economics has aspired to.

The agent-based approach should not and cannot replace the standard analytical economic approach. Rather, these two methodologies support and complement each other: When an analytical model is developed, it should become standard practice to examine the robustness of the model’s results with agent-based simulations. Similarly, when results emerge from agent-based simulation, one should try to understand their origin and their generality, not only by running many simulations but also by trying to capture the essence of the results in a simplified analytical setting (if possible).

Although the first steps in economic agent-based simulations were made decades ago, economics has been slow and cautious to adopt this new methodology. Only in recent years has this field begun to bloom. It is my belief and hope that the agent-based approach will prove as fruitful in economics as it has been in so many other branches of science.

Bibliography

Primary Literature

- Admati A, Pfleiderer P (1988) A theory of intraday patterns: volume and price variability. *Rev Financ Stud* 1:3–40
- Arthur WB (1994) Inductive reasoning and bounded rationality (The El Farol problem). *Am Econ Rev* 84:406–411
- Arthur WB, Holland JH, LeBaron B, Palmer RG, Tayler P (1997) Asset pricing under endogenous expectations in an artificial stock market. In: Arthur WB, Durlauf S, Lane D (eds) *The economy as an evolving complex system II*. Addison-Wesley, Redwood City
- Brock WA, Hommes CA (1998) Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *J Econ Dyn Control* 22:1235–1274
- Egenter E, Lux T, Stauffer D (1999) Finite size effects in Monte Carlo simulations of two stock market models. *Phys A* 268:250–256
- Epstein JM, Axtell RL (1996) Complex adaptive systems. In: *Growing artificial societies: social science from the bottom up*. MIT Press, Washington, DC
- Fama E, French K (1988) Permanent and temporary components of stock prices. *J Polit Econ* 96:246–273
- Friend I, Blume ME (1975) The demand for risky assets. *Am Econ Rev* 65:900–922
- Gordon J, Paradis GE, Rorke CH (1972) Experimental evidence on alternative portfolio decision rules. *Am Econ Rev* 62(1):107–118
- Grossman S, Stiglitz J (1980) On the impossibility of informationally efficient markets. *Am Econ Rev* 70:393–408
- Hellthaler T (1995) The influence of investor number on a microscopic market. *Int J Mod Phys C* 6:845–852
- Hommes CH (2002) Modeling the stylized facts in finance through simple nonlinear adaptive systems. *Proc Natl Acad Sci U S A* 99:7221–7228
- Jegadeesh N, Titman S (1993) Returns to buying winners and selling losers: implications for stock market efficiency. *J Finance* 48:65–91
- Karpoff J (1987) The relationship between price changes and trading volume: a survey. *J Finance Quant Anal* 22:109–126
- Kim GW, Markowitz HM (1989) Investment rules, margin, and market volatility. *J Portf Manage* 16:45–52
- Kirman AP (1992) Whom or what does the representative agent represent? *J Econ Perspect* 6:117–136
- Kohl R (1997) The influence of the number of different stocks on the Levy, Levy Solomon model. *Int J Mod Phys C* 8:1309–1316
- Kroll Y, Levy H, Rapoport A (1988) Experimental tests of the separation theorem and the capital asset pricing model. *Am Econ Rev* 78:500–519
- LeBaron B (2000) Agent-based computational finance: suggested readings and early research. *J Econ Dyn Control* 24:679–702
- Levy H (1994) Absolute and relative risk aversion: an experimental study. *J Risk Uncertain* 8:289–307
- Levy H, Lim KC (1998) The economic significance of the cross-sectional autoregressive model: further analysis. *Rev Quant Finance Account* 11:37–51
- Levy M, Levy H (1996) The danger of assuming homogeneous expectations. *Finance Analyst J* 52:65–70
- Levy M, Levy H, Solomon S (1994) A microscopic model of the stock market: cycles, booms, and crashes. *Econom Lett* 45:103–111

- Levy M, Levy H, Solomon S (2000) Microscopic simulation of financial markets. Academic, San Diego
- Levy M, Persky N, Solomon S (1996) The complex dyn of a simple stock market model. *Int J High Speed Comput* 8:93–113
- Lux T (1995) Herd behaviour, bubbles and crashes. *Econ J* 105:881
- Lux T (1998) The socio-economic dynamics of speculative bubbles: interacting agents, chaos, and the fat tails of returns distributions. *J Econ Behav Organ* 33:143–165
- Lux T, Marchesi M (1999) Volatility clustering in financial markets: a micro-simulation of interacting agents. *Nature* 397:498
- Orcutt GH, Caldwell SB, Wertheimer R (1976) Policy exploration through microanalytic simulation. The Urban Institute, Washington, DC
- Palmer RG, Arthur WB, Holland JH, LeBaron B, Tayler P (1994) Artificial economic life: a simple model of a stock market. *Phys D* 75:264–274
- Poterba JM, Summers LH (1988) Mean reversion in stock returns: evidence and implications. *J Finance Econ* 22:27–59
- Samanidou E, Zschischang E, Stauffer D, Lux T (2007) Agent-based models of financial markets. *Rep Prog Phys* 70:409–450
- Samuelson PA (1989) The judgement of economic science on rational portfolio management: timing and long horizon effects. *J Portf Manage* 16:4–12
- Samuelson PA (1994) The long term case for equities and how it can be oversold. *J Portf Manage* 21:15–24
- Sargent T (1993) Bounded rationality and macroeconomics. Oxford University Press, Oxford
- Schelling TC (1978) Micro motives and macro behavior. Norton, New York
- Shiller RJ (1981) Do stock prices move too much to be justified by subsequent changes in dividends? *Am Econ Rev* 71:421–436
- Stauffer D, de Oliveira PMC, Bernardes AT (1999) Monte Carlo simulation of volatility correlation in microscopic market model. *Int J Theor Appl Finance* 2:83–94
- Tesfatsion L (2001) Special issue on agent-based computational economics. *J Econ Dyn Control* 25:281–293
- Tesfatsion L (2002) Agent-based computational economics: growing economies from the bottom up. *Artif Life* 8:55–82
- Thaler R (ed) (1993) Advances in behavioral finance. Russel Sage Foundation, New York
- Thaler R (1994) Quasi rational economics. Russel Sage Foundation, New York
- Tversky A, Kahneman D (1981) The framing of decisions and the psychology of choice. *Science* 211:453–480
- Tversky A, Kahneman D (1986) Rational choice and the framing of decision. *J Bus* 59(4):251–278
- Tversky A, Kahneman D (1992) Advances in prospect theory: cumulative representation of uncertainty. *J Risk Uncertain* 5:297–323

Books and Reviews

- Anderson PW, Arrow J, Pines D (eds) (1988) The economy as an evolving complex system. Addison-Wesley, Redwood City
- Axelrod R (1997) The complexity of cooperation: agent-based models of conflict and cooperation. Princeton University Press, Princeton
- Moss de Oliveira S, de Oliveira H, Stauffer D (1999) Evolution, money, war and computers. BG Teubner, Stuttgart/Leipzig
- Solomon S (1995) The microscopic representation of complex macroscopic phenomena. In: Stauffer D (ed) *Annu Rev Comput Phys II*. World Scientific, Singapore

Cellular Automaton Modeling of Tumor Invasion

Haralambos Hatzikirou^{1,2}, Georg Breier³ and Andreas Deutsch¹

¹Center for Information Services and High Performance Computing, Technische Universität Dresden, Dresden, Germany

²Helmholtz Centre for Infection Research, Department Systems Immunology, Braunschweig, Germany

³Division of Medical Biology, Medical Faculty Carl Gustav Carus, Technische Universität Dresden, Dresden, Germany

Article Outline

Glossary

Definition of the Subject

Introduction

Cellular Automata

Models of Tumor Invasion

Invasive Tumor Morphology

Effects of Directed Cell Motion

Spatial Structure of Invasive Tumors

Tumor Cell Migration and the Influence of the Extracellular Matrix

The Role of Cell-Cell and Cell-ECM Adhesion

Cellular Mechanisms of Glioma Cell Migration

Effects of Fiber Tracts on Glioma Invasion

Effect of Heterogeneous Environments on Tumor Cell Migration

Metabolism and Acidosis

Emergence of Tumor Invasion

Influence of Metabolic Changes

The Game of Invasion

Discussion

Bibliography

Glossary

Cadherins Important class of transmembrane proteins. They play a significant role in cell-

cell adhesion, ensuring that cells within tissues are bound together.

Chemotaxis Motion response to chemical concentration gradients of a diffusive chemical substance.

Extracellular matrix (ECM) Components that are extracellular and composed of secreted fibrous proteins (e.g., collagen) and gel-like polysaccharides (e.g., glycosaminoglycans) binding cells and tissues together.

Fiber tracts Bundle of nerve fibers having a common origin, termination, and function within the spinal cord and brain.

Haptotaxis Directed motion of cells along adhesion gradients of fixed substrates in the ECM, such as integrins.

Slime trail motion Cells secrete a nondiffusive substance; concentration gradients of the substance allow the cells to migrate toward already explored paths.

Somatic evolution Darwinian-type evolution that occurs on soma (as opposed to germ) cells and characterizes cancer progression (Bodmer 1997).

Definition of the Subject

Cancer cells acquire characteristic traits in a step-wise manner during carcinogenesis. Some of these traits are autonomous growth, induction of angiogenesis, invasion, and metastasis. In this entry, the focus is on one of the late stages of tumor progression: tumor invasion. Tumor invasion has been recognized as a complex system, since its behavior emerges from the combined effect of tumor cell-cell and cell-microenvironment interactions. Cellular automata (CA) provide simple models of self-organizing complex systems in which collective behavior can emerge out of an ensemble of many interacting “simple” components. Cellular automata have also been used to gain a deeper insight in tumor invasion dynamics. In this entry, we briefly

introduce cellular automata as models of tumor invasion, and we critically review the most prominent CA models of tumor invasion.


Introduction

Cancer describes a group of genetic and epigenetic diseases characterized by uncontrolled growth of cells, leading to a variety of pathological consequences and frequently death. Cancer has long been recognized as an evolutionary disease (Nowell 1976). Cancer progression can be depicted as a sequence of traits or phenotypes that cells have to acquire if a neoplasm (benign tumor) is to become an invasive and malignant cancer. A phenotype refers to any kind of observed morphology, function, or behavior of a living cell. Hanahan and Weinberg (2000, 2011) have identified six cancer cell phenotypes: unlimited proliferative potential, environmental independence for growth, evasion of apoptosis, angiogenesis, invasion, and metastasis (Fig. 1).

In this entry, we focus on the invasive phase of tumor growth. Invasion is the main feature that allows a tumor to be characterized as malignant. The progression of a benign tumor and delimited growth to a tumor that is invasive and potentially metastatic is the major cause of poor clinical outcome in cancer patients, in terms of therapy and prognosis. Understanding tumor invasion could potentially lead to the design of novel therapeutical strategies. However, despite the immense amounts of funds invested in cancer research, the intracellular and extracellular dynamics that govern tumor invasiveness in vivo remain poorly understood.

Biomedically, invasion involves the following tumor cell *processes*:

- Tumor cell migration, which is a result of downregulation of certain cadherins, that is, loss of cell-cell adhesion (Breier et al. 2014).
- Tumor cell-extracellular matrix (ECM) interactions, such as cell-ECM adhesion, and ECM degradation or remodeling, by means of proteolysis. These processes allow for the penetration of the migrating tumor cells into host



Cellular Automaton Modeling of Tumor Invasion, Fig. 1 Hanahan and Weinberg have identified six possible types of cancer cell phenotypes: unlimited proliferative potential, environmental independence for growth, evasion of apoptosis, angiogenesis, invasion, and metastasis. (Reprinted from Hanahan and Weinberg 2000, with permission from the authors)

tissue barriers, such as basement and interstitial stroma (Friedl 2004).

- Tumor cell proliferation.

Tumor invasion facilitates the emergence of metastases, i.e., the spread of cancer cells to another part of the body and the formation of secondary tumors. Tumor invasion comprises a central aspect in cancer progression. However, invasive phenomena occur not only in pathological cases of malignant tumors but also during normal morphogenesis and wound healing.

Cancer research has been directed toward the understanding of tumor invasion dynamics and its implications in treatment design. In particular, research concentrates along the following *problems*.

Invasive tumor morphology	A wealth of empirical evidence links disease progression with tumor morphology (Sanga et al. 2007). The tumor
---------------------------	---

(continued)

	morphology can indicate the degree of a tumor's malignancy. In particular, it is experimentally and clinically observed that a morphological instability is related to invasive solid tumors, producing fingerlike spatial patterns. The question is which molecular and cellular mechanisms are responsible for this spatial pattern formation
Cell migration and influence of the ECM	Important aspects of invading tumors are cell motion and the effect of the surrounding environment, especially the ECM (Friedl 2004)
Metabolism and acidosis	The multistep process of carcinogenesis is often described by somatic evolution, wherein phenotypic properties are retained or lost depending on their contribution to the individual tumor cell survival and reproductive potential. One of the most prominent phenotypic changes involves the anaerobic glucose metabolism (glycolysis). A side product of this metabolic activity is the production of H^+ ions that increase the pH of the tumor's microenvironment (acidosis). This gives rise to the questions: (i) Why does tumor evolution lead to this kind of metabolism which is energetically deficient in comparison with the aerobic one? (ii) What are the advantages for the tumor? (iii) How do glycolytic tumor cells influence tumor invasion?
Emergence of invasion	Typically, tumor invasion appears during the late stages of carcinogenesis. Of ultimate importance is the question: What are the mechanisms and the environmental conditions that trigger the progression from benign neoplasms to malignant invasive tumors?
Robustness	There are several questions concerning the stability and the resistance of tumor invasion such as (i) Why are malignant tumors so robust (resistant) to

(continued)

	perturbations (i.e., therapies)? (ii) Is it possible to design intelligent therapies (at the cellular level) that disturb the tumor's robustness? (iii) How can we investigate the tumor's robustness?
--	--

Mathematical modeling and analysis provide invaluable tools toward answering the above questions. Tumor invasion involves processes which occur at different spatiotemporal scales, including processes at the subcellular, cellular, and tissue level. Mathematical models allow description and linking of these levels. One can distinguish *molecular*, *cellular*, and *tissue* scales, respectively (Hatzikirou et al. 2005; Preziosi 2003):

- The molecular scale refers to phenomena at the subcellular level and concentrates on molecular interactions and resulting phenomena, such as alterations of signaling cascades and cell cycle control, gene mutations, etc. In tumor invasion, the downregulation of cadherins provides an example of a molecular process.
- The cellular scale refers to cellular interactions and therefore to the most prominent dynamics of cell populations, e.g., adhesion, contact inhibition, chemotaxis, etc.
- The tissue scale focuses on tissue level processes taking into account macroscopic quantities, such as volumes, flows, etc. Continuum phenomena include cell convection and diffusion of nutrients and chemical factors, mechanical stress, and the diffusion of metastases.

For example, genetic alterations may lead to invasive cells (molecular scale) that are able to migrate (cellular scale) and interact with diffusible or nondiffusible signals (tissue scale). Models that deal with phenomena at multiple scales are called multi-scaled.

Meanwhile, a variety of mathematical models have been proposed to analyze different aspects of tumor invasion. Deterministic macroscopic models are used to model the spatiotemporal growth of tumors, usually assuming that tumor invasion is a wave propagation phenomenon

(Marchant et al. 2000; Perumpanani et al. 1996, 1999; Sherratt and Nowak 1992; Sherratt and Chaplain 2001). Computational investigations of the invasiveness of glioma tumors illustrate that the ratio of tumor growth and spatial anisotropy in cell motility can quantify the degree of tumor invasiveness (Swanson et al. 2002; Jbabdi et al. 2005). While these models are able to capture the tumor structure at the tissue level, they fail to describe the tumor at the cellular and the subcellular levels. Meanwhile, multi-scale approaches attempt to describe and predict invasive tumor morphologies, growth, and phenotypical heterogeneity (Anderson et al. 2006; Frieboes et al. 2007, see also Alfonso et al. 2017).

Cellular automata (CA), and more generally cell-based models, provide an alternative modeling approach, where a microscale investigation is allowed through a stochastic description of the dynamics at the cellular level (Deutsch and Dormann 2018). In particular, CA define an appropriate modeling framework for tumor invasion since they allow for the following:

- CA rules can mimic the processes at the cellular level. This fact allows for the modeling of an abundance of experimental data that refer to cellular and subcellular processes related to tumor invasion.
- The discrete nature of CA can be exploited for investigations of the boundary layer of a tumor de Franciscis et al. (2011). Bru et al. (2003) have analyzed the fractal properties of tumor surfaces (calculated by means of fractal scaling analysis) which can be compared with corresponding CA simulations to gain a better understanding of the tumor phenomenon. In addition, the discrete structure of CA facilitates the implementation of complicated environments without any of the computational problems characterizing the simulation of continuous models.
- Motion of tumor cells through heterogeneous media (e.g., ECM) involves phenomena at various spatial and temporal scales (Lesne 2007). These cannot be captured in a purely macroscopic modeling approach. Alternatively, discrete microscopic models, such as CA, can incorporate different spatiotemporal scales,

and they are well suited for simulating such phenomena.

- CA are paradigms of parallelizable algorithms. This fact makes them computationally efficient.

In the following section, we provide a definition of CA. In section “[Models of Tumor Invasion](#),” we review the existing CA models for central processes of tumor invasion. Finally, in the discussion, we critically discuss the use of CA in tumor invasion modeling, and we identify future research questions related to tumor invasion.

Cellular Automata

The notion of a *cellular automaton* originated in the works of John von Neumann (1903–1957) and Stanislaw Ulam (1909–1984). Cellular automata may be viewed as simple models of self-organizing complex systems in which collective behavior can emerge out of an ensemble of many interacting “simple” components. In complex systems, even if the basic and local interactions are perfectly known, it is possible that the global behavior obeys new laws that cannot be obviously extrapolated from the individual properties, as if the whole is more than the sum of the parts. This property makes cellular automata a very interesting approach to model complex systems in physics, chemistry, and biology (examples are introduced in Deutsch and Dormann 2018 ; Chopard et al. 2002). CA can be defined as a 4-tuple $\mathcal{L}, \mathcal{S}, \mathcal{N}, \mathcal{F}$, where:

- $\mathcal{L}$ is a finite or infinite regular *lattice* of nodes (discrete space)
- $\mathcal{S}$ is a finite set of *states* (discrete states); each cell $i \in \mathcal{L}$ is assigned a state $S \in \mathcal{S}$
- $\mathcal{N}$ is a finite set of *neighbors*
- $\mathcal{F}$ is a deterministic or probabilistic map

$$\mathcal{F} : \mathcal{S}^{|\mathcal{N}|} \rightarrow \mathcal{S}$$

$$\{\mathcal{S}_i\}_{i \in \mathcal{N}} \rightarrow \mathcal{S},$$

which assigns a new state to a node depending on the state of all its neighbors indicated by $\mathcal{N}$ (*local rule*).

The evolution of CA is defined by applying the function $\mathcal{F}$ synchronously to all nodes of the lattice $\mathcal{L}$ (*homogeneity* in space and time).

The above features can be extended, giving rise to several variants of the classical CA notion (Moreira and Deutsch 2002). Some of these are:

Asynchronous CA	In such CA, the restriction of simultaneous update of all the nodes is revoked, allowing for asynchronous update
Nonhomogeneous CA	This variation allows the transition rules to depend on node position. Agent-based models are “relatives” of CA that lost the homogeneity property, i.e., each individual particle may have its own set of rules
Coupled-map lattices	In this case the constraint of discrete state space is withdrawn, i.e., the state variables are assumed to be continuous. An important type of coupled map lattices is the so-called Lattice Boltzmann model (Succi 2001)
Structurally dynamic CA	In these systems, the underlying lattice is no longer a passive static object but becomes a dynamic component. Therefore, the lattice structure evolves depending on the values of the node’s state variables

An important class of cellular automaton models is lattice-gas cellular automata (LGCA). This CA model can describe discrete individuals interacting stochastically and moving in space. LGCA models were introduced to simulate aspects of fluid dynamics (Frisch et al. 1986) but have also been used successfully to investigate collective cell migration, biological pattern formation, and the growth, invasion and progression, of tumors (Böttger et al. 2012, 2015; Bussemaker et al. 1997; Chopard et al. 2010; de Franciscis 2011; Deutsch 1995, 2000; Dormann and Deutsch 2002; Dormann et al. 2001; Hatzikirou et al. 2015; Mente et al. 2012; Syga et al. 2019; Tektonidis et al. 2011; Mente et al. 2010; Buder et al. 2015, 2019; Dirkse et al. 2019; Alfonso et al. 2016, 2017; Reher et al. 2017; Talkenberger et al. 2017). LGCA models are cell-based and computationally efficient and allow to integrate statistical

and biophysical models for different levels of biological knowledge (Deutsch and Lawniczak 1999; Hatzikirou et al. 2010; Mente et al. 2012; Nava-Sedeño et al. 2017a, b, 2020a, b).

Models of Tumor Invasion

This section reviews the existing cellular automata models of tumor invasion. Categorizing these models is a nontrivial task. Moreover, existing CA models describe tumor invasion at more than one scale (subcellular, cellular, and tissue). In this entry, we distinguish models that analyze (i) the invasive morphology, (ii) tumor cell migration and the influence of the ECM, (iii) metabolism and acidosis, and (iv) the emergence of tumor invasion.

Invasive Tumor Morphology

The tumor morphology arising from the spatial pattern formation of the tumor cell population has been recognized as a very important aspect of tumor growth. Several researchers have attempted to reveal the mechanisms of spatial pattern formation of invasive tumors. Here, we present the most representative CA models for the invasive tumor morphology.

Effects of Directed Cell Motion

Sander and Deisboeck (2002) developed a CA model to investigate the branching morphology of invasive brain tumors. In the model tumor cell motion is influenced by two key processes: (i) chemotaxis and (ii) “slime trail following.” A typical example of a slime trail following mechanism is found in the motion of certain myxobacteria (Wolgemuth et al. 2002).

The authors show that the branching morphology of tumors can be explained as a result of chemotaxis and “slime trail following.” In particular, simulations reproduce the branching pattern formation observed in in vitro cultures of glioma

cells. However, the assumption of slime trail following has not been proven biologically as yet.

Spatial Structure of Invasive Tumors

Anderson (2005) and Anderson et al. (2006) proposed a model to examine the effects of tumor cell heterogeneity (at the genetic level) on the spatial morphology and to analyze the importance of cell-cell and cell-ECM adhesion. The model assumes a nondiffusible, fixed configuration of ECM. The extracellular matrix can be degraded by diffusible enzymes, such as metalloproteinases, produced by tumor cells. Moreover, cells are allowed to mutate and evolve their phenotype from proliferative to invasive. Finally, an oxygen concentration field plays the role of nutrients in the model.


Simulations of the model show that (i) the ECM heterogeneity is mainly responsible for the tumor branching morphology (Fig. 2), (ii) cell-cell adhesion plays an important role only in the early stages of tumor development, (iii) invasive tumor cells are located at the boundary of the tumor, and (iv) the tumor is a phenotypically heterogeneous object.

Tumor Cell Migration and the Influence of the Extracellular Matrix

Cell migration and cell-ECM interactions are two of the most crucial invasion-related processes. Cellular automata provide an appropriate framework to model and analyze the effect of cell motility and cell-environment interactions of tumor cell migration.

The Role of Cell-Cell and Cell-ECM Adhesion

Turner and Sherratt (2002) proposed a cellular Potts model (Graner and Glazier 1992) to investigate how cell-cell and cell-ECM adhesion influence the tumor invasion depth and tumor morphology. A cellular Potts model can be viewed as an extension of the CA idea allowing to analyze phenomena that take into account specific cell shapes. Cells are assumed to move according to intercellular adhesive interactions and haptotactical gradients. Moreover, cells are allowed to proliferate, while mitotic probabilities depend on the strength of the adhesive interaction. Finally, cells are assumed to secrete proteolytic enzymes that degrade the ECM.



Cellular Automaton Modeling of Tumor Invasion, Fig. 2 *Left:* Microscopy image of a multicellular tumor spheroid, exhibiting an extensive branching system that rapidly expands into the surrounding extracellular matrix gel. These branches consist of multiple invasive cells.

(Reprinted from Habib et al. 2003 with permission.) *Right:* Simulation of Anderson's model (Anderson et al. 2006) reproducing the experimentally observed morphology of invasive tumors

The authors show that adhesive dynamics can explain the “fingering” patterns observed in their simulations. Moreover, the authors demonstrate that the width of the invasion zone depends less on cell-cell adhesion and more on cell-ECM adhesion facilitated by haptotaxis and proteolysis.

Cellular Mechanisms of Glioma Cell Migration

In the work of Aubert et al. (2006), a CA model is introduced that allows for the investigation of tumor cell migration, based on experimentally observed density profiles of glioma cell cultures. The goal is to identify the mechanisms of tumor (glioma) cell motion, which play a crucial role in tumor invasion. The authors do not consider proliferation of tumor cells. Only the influences of tumor cell migration and intercellular interactions are studied. The authors introduce and test two distinct cell mechanisms: (i) cell-cell adhesion and (ii) a kind of “inertia” in cell motion, i.e., the cells tend to maintain the direction of their motion.

The authors carefully scale the model according to the experimental setup and calibrate the corresponding model parameters. The simulation results indicate that cell-cell adhesion can explain the experimental results. It is concluded that cell-cell adhesion is an important process in glioma cell migration.

Effects of Fiber Tracts on Glioma Invasion

Wurzel et al. (2005) model glioma tumor invasion with a lattice-gas cellular automaton (LGCA) (Deutsch and Dormann 2018). The authors address the question of how fiber tracts found in the brain’s white matter influence the spatiotemporal evolution and the invading front morphology of glioma tumors. Cells are assumed to move, proliferate, and undergo apoptosis according to corresponding stochastic processes. Fiber tracts are represented as a local gradient field that enhances cell motion in a specific direction.

The authors develop and analyze different scenarios of fiber tract influence. A gradient field may increase the speed of the invading tumor front. For high field intensities, the model predicts the formation of cancer islets at distances away from the main tumor bulk. The simulated invasion patterns qualitatively resemble clinical observations.

Effect of Heterogeneous Environments on Tumor Cell Migration


Hatzikirou and Deutsch (2008) developed an LGCA model to investigate the influence of heterogeneous environments on tumor cell dispersal. This model is a simplified version of Wurzel et al. (2005) which facilitates the mathematical analysis. In this study no proliferation or death of cells is considered. The authors distinguish two kinds of cell-ECM interactions: (i) cell-ECM adhesion leading to haptotactical motion along integrin concentration gradients (environment with directional information) and (ii) contact guidance that promotes the alignment along ECM pores or fibers as seen in Fig. 3 (environment with orientational information).

In this study, the impact of both types of cell-ECM interaction on tumor cell motion is investigated. In particular, macroscopic dispersal measures (such as mean cell flux) depending on cellular and environmental parameters are calculated. Accordingly, the models allow for prediction of cell motion in different environments.

Metabolism and Acidosis

In the course of cancer progression, tumor cells undergo several phenotypic changes in terms of motility, metabolism, and proliferative rates. In particular, it is important to analyze the effect of the anaerobic metabolism of tumor cells and the acidification of the environment (as a side product of glycolysis) on tumor invasion (Fig. 4).

Patel et al. (2001) proposed a model of tumor growth to examine the roles of native tissue vascularity and anaerobic metabolism on the growth and invasion efficacy of tumors. The model



Cellular Automaton Modeling of Tumor Invasion, Fig. 3 The effect of the brain's fiber tracts on glioma growth. (a) A simulation is shown without taking into account the influence of fiber tracts. (b) The fiber tracts in

the brain strongly drive the evolution of the tumor growth. (c, d) Figures display a close-up of the tumor area of the top (a, b) simulations. (Reprinted from Hatzikirou and Deutsch 2008)

assumes a vascularized host tissue. Anaerobic metabolism involves the consumption of glucose and the production of H^+ ions, leading to the acidification of the local tissue. The vascular network allows for the absorption of H^+ ions. Cells are assumed to be proliferative and non-motile. The pH level, i.e., the H^+ concentration, and the glucose concentration determine the survival and death of the cells.


Simulations of the model show that (i) high tumor H^+ ion production favors tumor invasion by the acidification of the neighboring host tissue

and (ii) there is an optimal density of microvessels that maximizes tumor growth and invasion, by minimizing the acidification effects on tumor cell proliferation (absorption of H^+ ions) and maximizing the negative effect of H^+ ions on the neighboring tissue.

Emergence of Tumor Invasion

Several models have been proposed that concentrate on the evolutionary dynamics of tumors

(Fig. 5). The main goal of these models is to understand under which environmental conditions particular phenotypes appear. Here, we review those models that focus on the mechanisms that allow the emergence of invasive behavior.




Cellular Automaton Modeling of Tumor Invasion, Fig. 4 Typically, tumors exhibit abnormal levels of glucose metabolism. Positron emission imaging (PET) techniques localize the regions of abnormal glycolytic activity and identify the tumor locus

Influence of Metabolic Changes

Smallbone et al. (2007) developed an evolutionary CA model to investigate the cell-microenvironment interactions that mediate somatic evolution of cancer cells. In particular, the authors investigate the sequence of tumor phenotypes that ultimately leads to invasive behavior. The model considers three phenotypes: (i) the hyperplastic phenotype that allows growth away from the basement membrane, (ii) the glycolytic phenotype that allows anaerobic metabolism (the “fuel” is glucose), and (iii) the acid-resistant phenotype that enables the cell to survive in low pH. Cells are allowed to proliferate, die, or adapt (change their phenotype). No cell motion is explicitly considered.

The model predicts three phases of somatic evolution: (i) Initially, cell survival and proliferation are dependent on the oxygen concentration. (ii) When the oxygen becomes scarce, the glycolytic phenotype confers a significant proliferative advantage. (iii) The side products of glycolysis, e.g., galactic acid, increase the microenvironmental pH and promote the selection of acid-resistant phenotypes. The latter cell type is able to invade the neighboring tissue since it takes advantage of

Cellular Automaton Modeling of Tumor Invasion, Fig. 5 The evolving microenvironment of breast cancer. The multiple stages of breast carcinogenesis are shown progressing from *left to right*, along with histological representations of these stages. As indicated, the preinvasive stages occur in an avascular environment, whereas cancer cells have direct access to vasculature following invasion. (Reprinted from Gillies and Gatenby 2007)



the death of host cells, due to acidification, and proliferates using the available free space.

The Game of Invasion

Basanta et al. (2008) have developed a game theory-inspired CA that addresses the question of how invasive behavior emerges during tumor progression (see also Basanta et al. 2008; Hummert et al. 2014). The authors study the circumstances under which mutations that confer increased motility to cells can spread through a tumor composed of rapidly proliferating cells. The model assumes the existence of only two phenotypes: “proliferative” (high division rate and no motility) and “migratory” (low division rate and high motility). Mutations are allowed for by the random change of phenotypes. Nutrients are assumed to be uniformly distributed over the lattice.

Simulations show that low-nutrient conditions confer a reproductive advantage to motile cells over the proliferative ones. The model suggests novel ideas for therapeutic strategies, e.g., by increasing the oxygen supply around the tumor to favor the reproduction of proliferative cells over the migrating ones. This is not necessarily a therapy since there are benign tumors that are life-threatening even if they do not become invasive. Despite that, in most cases a growing but non-aggressive tumor will have a much better prognosis than a smaller but invasive one.

Discussion

In this entry, we have focused on one of the most important aspects of cancer progression: tumor invasion. The main processes involved in tumor invasion are related to tumor cell migration and cell-ECM interactions, especially ECM degradation/remodeling and tumor cell proliferation. These processes are evolving at different scales, e.g., cell-ECM adhesion is the response of tumor cells to ECM integrins (molecular level) leading to a haptotactical cell motion (cellular level) and influencing the tumor morphology (macroscopic level). Therefore, in order to understand tumor

invasion dynamics, it is important to use mathematical tools that allow for modeling subcellular or cellular processes and to analyze the emergent macroscopic behavior. Individual-based models, especially CA, are well suited for this task. Moreover, some types of CA models, such as lattice-gas cellular automata (Deutsch and Dormann 2018; Hatzikirou and Deutsch 2008), facilitate analytical investigations allowing for deeper insight into the modeled phenomena.

In this entry, we reviewed the existing CA models of tumor invasion. The presented models explore central aspects of tumor invasion. Some of the models are in good agreement with biomedical observations for in vitro and in vivo tumors. In the following, we list the most interesting biological insights that can be gained from the reviewed models:

- The significance of hypoxia in the process of tumor progression: Activation of glycolysis and acidification of the host tissue facilitate tumor invasion. Low-nutrient conditions, such as hypoxia, may trigger invasive behaviors.
- Cell-cell adhesion: It is evident that intercellular adhesion has a great impact in the early stages of tumor growth. However, in tumor invasion the role of cell-cell adhesion is minor, since mainly the cell-ECM interactions appear to dictate the tumor cell behavior.
- Cell-ECM adhesion: This is an important process for tumor invasion. In particular, the heterogeneous structure of the ECM strongly influences the spatial morphology of invasive tumors.

Mathematical modeling offers potentially significant insight into tumor invasion. Several crucial questions have not been adequately addressed so far by modeling efforts.

Branching morphology	Several mechanisms have been proposed that lead to branching patterns, e.g., diffusion-limited aggregation, the interplay of cell-cell and cell-ECM adhesion, as well as chemotaxis or slime trail following motion. However, biologists and modelers have not yet identified a
----------------------	---

(continued)

	unique mechanism that drives the branching morphology of invasive tumors
Go or grow	The mechanisms of invasive tumor cell migration are still not understood. Fedotov and Iomin (2007) have analyzed the effect of a postulated migration/proliferation dichotomy on cell migration
Emergence	Concerning the emergence of invasion in tumor progression, little is known. Mechanisms related to tumor cell motion and other cell processes, such as proliferation (migration/proliferation dichotomy), may play an important role for the dominance of invasive phenotypes (Basanta et al. 2008; Hatzikirou et al. 2012)
Angiogenesis	Another open issue is the influence of angiogenesis and vasculogenesis on tumor invasion. Despite significant efforts to describe the mechanisms of angio- and vasculogenesis, little is known about the effect of these processes on tumor invasion (Patel et al. 2001; Alfonso et al. 2016)
Robustness	The identification of cellular mechanisms that are responsible for tumor robustness remains a significant challenge

Finally, for clinical purposes, future models should be able to provide accurate and quantitative predictions. Simplified models considering only the essential ingredients for tumor growth, and especially tumor invasion, but validated with actual clinical data may be helpful in this regard. We sincerely hope that a more profound knowledge of important tumor characteristics, such as tumor invasion, will eventually lead to the design of more effective therapeutic strategies.

Acknowledgments We are grateful for generous support throughout the years by the Centre for Information Services and High Performance Computing, Dresden University of Technology, Germany.

Bibliography

Alfonso JCL, Kohn-Luque A, Stylianopoulos T, Feuerhake F, Deutsch A, Hatzikirou H (2016) Why one-size-fits-all vaso-modulatory interventions fail to control glioma invasion: in silico insights Scientific Reports, 6:37283

Alfonso JCL, Talkenberger K, Seifert M, Klink B, Hawkins-Daarud A, Hatzikirou H, Swanson K, Deutsch A (2017) The biology and mathematical modelling of glioma invasion: a review. *J R Soc Interf* 14(136):20170490

Anderson ARA (2005) A hybrid model of solid tumour invasion: the importance of cell adhesion. *Math Med Biol* 22:163–186

Anderson A, Weaver A, Cummings P, Quaranta V (2006) Tumor morphology and phenotypics evolution driven by selective pressure from the microenvironment. *Cell* 127:905–915

Aubert M, Badoual M, Freol S, Christov C, Grammaticos B (2006) A cellular automaton model for the migration of glioma cells. *Phys Biol* 3:93–100

Basanta D, Hatzikirou H, Deutsch A (2008) The emergence of invasiveness in tumours: a game theoretic approach. *Eur Phys J B* 63:393–397

Basanta D, Simon M, Hatzikirou H, Deutsch A (2008) An evolutionary game theory perspective elucidates the role of glycolysis in tumour invasion. *Cell Prolif* 41:980

Bodmer W (1997) Somatic evolution of cancer cells. *J R Coll Physicians Lond* 31(1):82–89

Böttger K, Hatzikirou H, Chauviere A, Deutsch A (2012) Investigation of the migration/proliferation dichotomy and its impact on avascular glioma invasion. *Math Model Nat Phenom* 7:105–135

Böttger K, Hatzikirou H, Voss-Böhme A, Cavalcanti-Adam EA, Herrero MA, Deutsch A (2015) An emerging Allee effect is critical for tumor initiation and persistence. *PLOS Comput Biol* 11:1–14. <https://doi.org/10.1371/journal.pcbi.1004366>

Breier G, Grosser M, Rezaei M (2014) Endothelial cadherins in cancer. *Cell Tissue Res* 355:523–527

Bru A, Albertos S, Subiza JL, Lopez Garcia-Asenjo J, Bru I (2003) The universal dynamics of tumor growth. *Biophys J* 85:2948–2961

Buder T, Deutsch A, Klink B, Voss-Böhme A (2015) Model-based evaluation of spontaneous tumor regression in pilocytic astrocytoma *PLoS Comput Biol*, 11 (12):e1004662

Buder T, Deutsch A, Klink B, Voss-Böhme A (2019) Patterns of tumor progression predict small and tissue-specific tumor-originating niches *Front. Oncol.*, 8, 668

Bussemaker HJ, Deutsch A, Geigant E (1997) Mean-field analysis of a dynamical phase transition in a cellular automaton model for collective motion. *Phys Rev Lett* 78:5018–5021. <https://doi.org/10.1103/PhysRevLett.78.5018>

Chopard B, Dupuis A, Masselot A, Luthi P (2002) Cellular automata and lattice Boltzmann techniques: an approach to model and simulate complex systems. *Adv Compl Syst* 5(2):103–246

Chopard B, Ouared R, Deutsch A, Hatzikirou H, Wolf-Gladrow D (2010) Lattice-gas cellular automaton models for biology: from fluids to cells. *Acta Biotheor* 58:329–340

de Francis S, Hatzikirou H, Deutsch A (2011) Analysis of lattice-gas cellular automaton models for tumor growth by means of fractal scaling. *Acta Phys Pol B Proc Suppl* 4:167

- Deutsch A (1995) Towards analyzing complex swarming patterns in biological systems with the help of lattice-gas cellular automata. *J Biol Syst* 3:947–955
- Deutsch A (2000) A new mechanism of aggregation in a lattice-gas cellular automaton model. *Math Comput Model* 31:35–40
- Deutsch A, Dormann S (2018) Cellular automaton modeling of biological pattern formation. Birkhauser, Boston
- Deutsch A, Lawniczak AT (1999) Probabilistic lattice models of collective motion and aggregation: from individual to collective dynamics. *Math Biosci* 156:255–269
- Dirkse A, Golebiewska A, Buder T, Nazarov PV, Muller A, Poovathingal S, Brons NHC, Leite S, Sauvageot N, Sarkisjan D, Seyfrid M, Fritah S, Stieber D, Michelucci A, Hertel F, Herold-Mende C, Azuaje F, Skupin A, Bjerkvig R, Deutsch A, Voss-Böhme A, Niclou SP (2019) Stem cell-associated heterogeneity in Glioblastoma results from intrinsic tumor plasticity shaped by the microenvironment *Nature Communications*, 10(1):1787
- Dormann S, Deutsch A (2002) Modeling of self-organized avascular tumor growth with a hybrid cellular automaton. *Silico Bio* 2:393–406
- Dormann S, Deutsch A, Lawniczak AT (2001) Fourier analysis of Turing-like pattern formation in cellular automaton models. *Futur Gener Comput Syst* 17:901–909. [https://doi.org/10.1016/S0167-739X\(00\)00068-6](https://doi.org/10.1016/S0167-739X(00)00068-6)
- Fedotov S, Iomin A (2007) Migration and proliferation dichotomy in tumor-cell invasion. *Phys Rev Lett* 98:118101–118104
- Frieboes H, Lowengrub J, Wise S, Zheng X, Macklin P, Bearer E, Cristini V (2007) Computer simulation of glioma growth and morphology. *NeuroImage* 37(1):59–70
- Friedl P (2004) Prespecification and plasticity: shifting mechanisms of cell migration. *Curr Opin Cell Biol* 16(1):14–23
- Frisch U, Hasslacher B, Pomeau Y (1986) Lattice-gas automata for the Navier-Stokes equation. *Phys Rev Lett* 56:1505–1508
- Gillies RJ, Gatenby RA (2007) Hypoxia and adaptive landscapes in the evolution of carcinogenesis. *Cancer Metastasis Rev* 26:311–317
- Graner F, Glazier J (1992) Simulation of biological cell sorting using a two-dimensional extended Potts model. *Phys Rev Lett* 69:2013–2016
- Habib S, Molina-Paris C, Deisboeck TS (2003) Complex dynamics of tumors: modeling an emerging brain tumor system with coupled reaction-diffusion equations. *Phys A* 327:501–524
- Hanahan D, Weinberg R (2000) The hallmarks of cancer. *Cell* 100:57–70
- Hanahan D, Weinberg R (2011) Hallmarks of cancer. The next generation. *Cell* 144:646–674
- Hatzikirou H, Deutsch A (2008) Cellular automata as microscopic models of cell migration in heterogeneous environments. *Curr Top Dev Biol* 81:401–434
- Hatzikirou H, Deutsch A, Schaller C, Simon M, Swanson K (2005) Mathematical modelling of glioblastoma tumour development: a review. *Math Models Method Appl Sci* 15(11):1779–1794
- Hatzikirou H, Brusch L, Deutsch A (2010) From cellular automaton rules to a macroscopic mean-field description. *Acta Phys Pol B Proc Suppl* 3:399–416
- Hatzikirou H, Basanta B, Simon M, Schaller C, Deutsch A (2012) “Go or grow”: the key to the emergence of invasion in tumor progression? *Math Med Biol* 29(1):49–65
- Hatzikirou H, Böttger K, Deutsch A (2015) Model-based comparison of cell density-dependent cell migration strategies. *Math Model Nat Phenom* 10:94–107
- Hummert S, Bohl K, Basanta D, Deutsch A, Werner S, Theißen G, Schroeter A, Schuster S (2014) Evolutionary game theory: cells as players *Mol. BioSyst.*, 10, 3044–3065
- Jbabdi S, Mandonnet E, Duffau H, Capelle L, Swanson K, Pelegrini-Issac M, Guillemin R, Benali H (2005) Simulation of anisotropic growth of low-grade gliomas using diffusion tensor imaging. *Magn Reson Med* 54:616–624
- Lesne A (2007) Discrete vs continuous controversy in physics. *Math Struct Comput Sci* 17:185–223
- Marchant BP, Norbury J, Perumpanani AJ (2000) Traveling shock waves arising in a model of malignant invasion. *SIAM J Appl Math* 60(2):263–276
- Mente C, Prade I, Brusch L, Breier G, Deutsch A (2010) Parameter estimation with a novel gradient-based optimization method for biological lattice-gas cellular automaton models. *J Math Bio* 63:173–200
- Mente C, Prade I, Brusch L, Breier G, Deutsch A (2012) A lattice-gas cellular automaton model for in vitro sprouting angiogenesis. *Acta Phys Pol B* 5:99–115
- Moreira J, Deutsch A (2002) Cellular automaton models of tumour development: a critical review. *Adv Compl Syst* 5:1–21
- Nava-Sedeño JM, Hatzikirou H, Klages R, Deutsch A (2017a) Cellular automaton models for time-correlated random walks: derivation and analysis. *Sci Rep* 7:16952
- Nava-Sedeño JM, Hatzikirou H, Peruani F, Deutsch A (2017b) Extracting cellular automaton rules from physical Langevin equation models for single and collective cell migration. *J Math Biol* 75:1075–1100
- Nava-Sedeño JM, Voss-Böhme A, Hatzikirou H, Deutsch A, Peruani F (2020) Modeling collective cell motion: are on- and off-lattice models equivalent? *Roy. Soc. Open Sc*
- Nava-Sedeno JM, Hatzikirou H, Voss-Böhme A, Brusch L, Deutsch A, Peruani F (2020) Vectorial active matter on the lattice: emergence of polar condensates and nematic bands in an active zero-range process hal-02460291
- Nowell PC (1976) The clonal evolution of tumor cell populations. *Science* 194(4260):23–28
- Patel A, Gawlinski E, Lemieux S, Gatenby R (2001) Cellular automaton model of early tumor growth and invasion: the effects of native tissue vascularity and increased anaerobic tumor metabolism. *J Theor Biol* 213:315–331
- Perumpanani AJ, Sherratt JA, Norbury J, Byrne HM (1996) Biological inferences from a mathematical

- model of malignant invasion. *Invasion Metastasis* 16:209–221
- Perumpanani AJ, Sherratt JA, Norbury J, Byrne HM (1999) A two parameter family of travelling waves with a singular barrier arising from the modelling of extracellular matrix mediated cellular invasion. *Phys D* 126:145–159
- Preziosi L (ed) (2003) *Cancer modelling and simulation*. Chapman & Hall/CRC Press, Boca Raton
- Reher D, Klink B, Deutsch A, Voss-Böhme A (2017) Cell adhesion heterogeneity reinforces tumour cell dissemination: novel insights from a mathematical model *Biology Direct*, 12(1):18
- Sander LM, Deisboeck TS (2002) Growth patterns of microscopic brain tumours. *Phys Rev E* 66:051901
- Sanga S, Frieboes H, Zheng X, Gatenby R, Bearer E, Cristini V (2007) Predictive oncology: multidisciplinary, multi-scale in-silico modeling linking phenotype, morphology and growth. *NeuroImage* 37(1):120–134
- Sherratt JA, Chaplain MAJ (2001) A new mathematical model for avascular tumour growth. *J Math Biol* 43:291–312
- Sherratt JA, Nowak MA (1992) Oncogenes, anti-oncogenes and the immune response to cancer: a mathematical model. *Proc R Soc Lond B* 248:261–271
- Smallbone K, Gatenby R, Gillies R, Maini P, Gavaghan D (2007) Metabolic changes during carcinogenesis: potential impact on invasiveness. *J Theor Biol* 244:703–713
- Succi S (2001) *The lattice Boltzmann equation: for fluid dynamics and beyond*, series numerical mathematics and scientific computation. Oxford University Press, Oxford
- Swanson KR, Alvord EC, Murray J (2002) Quantifying efficacy of chemotherapy of brain tumors (gliomas) with homogeneous and heterogeneous drug delivery. *Acta Biotheor* 50:223–237
- Syga S, Nava-Sedeño JM, Brusch L, Deutsch A (2019) A lattice-gas cellular automaton model for discrete excitable media, chapter 15. In: Müller S, Tsuji K (eds) *Spirals and vortices*. Springer, Cham, pp 253–264, Springer
- Talkenberger K, Cavalcanti-Adam EA, Voss-Böhme A, Deutsch A (2017) Amoeboid-mesenchymal migration plasticity promotes invasion only in complex heterogeneous microenvironments *Scientific Reports*, 7:9237
- Tektonidis M, Tektonidis HH, Simon M, Schaller C, Deutsch A (2011) Identification of intrinsic in vitro cellular mechanisms for glioma invasion. *J Theor Bio* 287:131–147
- Turner S, Sherratt JA (2002) Intercellular adhesion and cancer invasion: a discrete simulation using the extended Potts model. *J Theor Biol* 216:85–100
- Wolgemuth CW, Hoiczyk E, Kaiser D, Oster GF (2002) How myxobacteria glide. *Curr Biol* 12(5):369–377
- Wurzel M, Schaller C, Simon M, Deutsch A (2005) Cancer cell invasion of normal brain tissue: guided by pre-pattern? *J Theor Med* 6(1):21–31

Agent-Based Modeling and Computer Languages

Michael J. North¹ and Charles M. Macal²

¹Argonne National Laboratory, Global Security Sciences Division, Argonne, IL, USA

²Center for Complex Adaptive Agent Systems Simulation (CAS2), Decision and Information Sciences Division, Argonne National Laboratory, Argonne, IL, USA

Article Outline

Glossary

Definition: Agent-Based Modeling and Computer Languages

Agent-Based Modeling

Types of Computer Languages

Requirements of Computer Languages for Agent-Based Modeling

Example Computer Languages Useful for Agent-Based Modeling

Future Directions

Bibliography

Keywords

Agent-based mode · Agent-based simulation · Computer language · Complex adaptive systems modeling

Glossary

Agent An agent is a self-directed component in an agent-based model.

Agent-Based Model An agent-based model is a simulation made up of a set of agents and an agent interaction environment.

Annotations Annotations are a Java feature for including metadata in compiled code.

Aspects Aspects are a way to implement dispersed but recurrent tasks in one location.

Attributes Attributes are a C# feature for including metadata in compiled code.

Bytecode Bytecode is compiled Java binary code.

C# C# (Archer 2001) is an object-oriented programming language that was developed and is maintained by Microsoft. C# is one of many languages that can be used to generate Microsoft.NET Framework code. This code is run using a “virtual machine” that potentially gives it a consistent execution environment on different computer platforms.

C++ C++ is a widely used object-oriented programming language that was created by Bjarne Stroustrup (Stroustrup 2008) at AT&T. C++ is widely used for both its object-oriented structure and its ability to be easily compiled into native machine code.

Class A class is the object-oriented inheritable binding of procedures and data that provides the basis for creating objects.

Common Intermediate Language Common Intermediate Language (CIL) is compiled binary code for the Microsoft.NET Framework. CIL was originally called Microsoft Intermediate Language (MSIL).

Computational Algebra Systems Computational Algebra Systems (CAS) are computational mathematics systems that calculate using symbolic expressions.

Computational Mathematics Systems Computational Mathematics Systems (CMS) are software programs that allow users to apply powerful mathematical algorithms to solve problems through a convenient and interactive user interface. CMS typically supply a wide range of built-in functions and algorithms.

Computer Language A computer language is a method of specifying directives for computers. Computer programming languages, or more simply programming languages, are an important category of computer languages.

Computer Programming Language Please see the entry for “Programming Language.”

Declarative Language According to Watson (1989) a “declarative language (or non-

procedural language) involves the specification of a set of rules defining the solution to the problem; it is then up to the computer to determine how to reach a solution consistent with the given rules.”

Design Pattern Design patterns form a “common vocabulary” describing tried-and-true solutions for commonly faced software design problems (Coplien 2001).

Domain-Specific Language Domain-specific languages (DSLs) are computer languages that are highly customized to support a well-defined application area or “domain.” DSLs commonly include a substantial number of keywords that are nouns and verbs in the area of application as well as overall structures and execution patterns that correspond closely with the application area.

Dynamic Method Invocation Dynamic method invocation, combined with reflection, is a Java and C# approach to higher-order programming.

Encapsulation Encapsulation is the containment of details inside a module.

Field A field is a piece of object-oriented data.

Function Pointers Function pointers are part of C++’s approach to higher-order programming. Runtime Type Identification is another component of this approach.

Functional Language According to Watson (1989) “in functional languages (sometimes called applicative languages) the fundamental operation is function application.”

Generics Generics are a Java and C# feature for generalizing classes.

Goto Statement A goto statement is an unconditional jump in a code execution flow.

Headless A headless program executes without the use of a graphical user interface or video monitor. This is generally done to rapidly execute models while logging results to files or databases.

Higher-Order Programming According to Reynolds (1998), higher-order programming involves the use of “procedures or labels. . . as data” such that they “can be used as arguments to procedures, as results of functions, or as values of assignable variables.”

Imperative Language According to Watson (1989) in imperative languages “there is a fundamental underlying dependence on the assignment operation and on variables implemented as computer memory locations, whose contents can be read and altered.”

Inheritance Inheritance is the ability of an object-oriented class to assume the methods and data of another class called the parent class.

Java Java (Foxwell 1999) is a widely used object-oriented programming language that was developed and is maintained by Oracle Corporation. Java is known for its widespread cross-platform availability on many different types of hardware and operating systems. This capability comes from Java’s use of a “virtual machine” that allows code to have a consistent execution environment on many different computer platforms.

Logic Programming Language According to Watson (1989) “in a logic programming language, the programmer needs only to supply the problem specification in some formal form, as it is the responsibility of the language system to infer a method of solution.”

Macro Language Macro languages are simple domain-specific languages that are used to write scripts for tools such as spreadsheets.

Mathematica *Mathematica* is a commercial software program for computational mathematics. Information on *Mathematica* can be found at <http://www.wolfram.com/>

MATLAB The MATrix LABoratory (MATLAB) is a commercial software program for computational mathematics. Information on MATLAB can be found at <http://www.mathworks.com/>

Method A method is an object-oriented procedure.

Methodological Individualism A reductionist approach to social science originally developed by Max Weber that focuses on the interaction of well-defined and separate individuals (Heath 2005). Alternative theories usually focus on more holistic views of interaction (Heath 2005).

Mobile Agents Mobile agents are lightweight software proxies that roam the World Wide Web and perform various functions programmed by their owners such as gathering information from Web sites.

- Module** According to Stevens et al. (1974), “the term module is used to refer to a set of one or more contiguous program statements having a name by which other parts of the system can invoke it and preferably having its own distinct set of variable names.”
- NetLogo** NetLogo (Wilensky 1999) is an agent-based modeling and simulation platform that uses a domain-specific language to define models. NetLogo models are built using a metaphor of turtles as agents and patches as environmental components (Wilensky 1999). NetLogo is Java based. NetLogo is free for use in education and research. More information on NetLogo and downloads can be found at <http://ccl.northwestern.edu/netlogo/>
- Non-procedural Language** Please see the entry for Declarative Language.
- Object** An object is the instantiation of a class to produce executable instances.
- Objective-C** Objective-C is an object-oriented language that extends the C language.
- Object-Oriented Language** Object-oriented languages are structured languages that have special features for binding data with procedures, inheritance, encapsulation, and polymorphism. Careful abstraction that avoids unnecessary details is an important design principle associated with the use of object-oriented languages.
- Observer** The observer is a NetLogo agent that has a view of an entire model. There is exactly one observer in every NetLogo model.
- ODD Protocol** Describes models using a three-part natural language approach: overview, concepts, and details (Grimm et al. 2006).
- Patch** A patch is a NetLogo agent with a fixed location on a master grid.
- Polymorphism** Polymorphism is the ability of an object-oriented class to respond to multiple related messages, often method calls with the same name but different parameters.
- Procedural Language** According to Watson (1989) “procedural languages...are those in which the action of the program is defined by a series of operations defined by the programmer.”
- Programming Language** A programming language is a computer language that allows any computable activity to be expressed.
- Record** A record is an independently addressable collection of data items.
- Reflection** Reflection, combined with dynamic method invocation, is a Java and C# approach to higher-order programming.
- ReLogo** An object-oriented Logo implementation in Repast Simphony.
- Repast** The Recursive Porous Agent Simulation Toolkit (Repast) is a free and open source family of agent-based modeling and simulation platforms (ROAD 2013). Information on Repast and free downloads can be found at <http://repast.sourceforge.net/>
- Repast Simphony** Repast Simphony is the member of the Repast Suite of free and open source agent-based modeling and simulation software (North et al. 2013). The Java-based Repast Simphony system includes advanced features for specifying, executing, and analyzing agent-based simulations.
- Runtime Type Identification** Runtime Type Identification (RTTI) is part of C++’s approach to higher-order programming. Function pointers are another component of this approach.
- Structured Language** Structured languages are languages that divide programs into separate modules, each of which has one controlled entry point, a limited number of exit points, and no internal jumps (Dijkstra 1968).
- Swarm** Swarm (Swarm Development Group 2013) is a free and open source agent-based modeling and simulation library maintained by the Swarm Development Group. The core Swarm system uses Objective-C. A Java-based “Java Swarm” wrapper for the Objective-C core is also available. Information on Swarm and free downloads can be found at <http://www.swarm.org/>
- Templates** Templates are a C++ feature for generalizing classes.
- Turtle** A turtle is a mobile NetLogo agent.
- Unified Modeling Language** The Unified Modeling Language (UML) is a predominantly visual approach to specifying the design of software (Object Management Group 2001, Object Management Group 2013) that consists of a variety of diagram types.

Unstructured Language Unstructured languages are languages that rely on step-by-step solutions such that the solutions can contain arbitrary jumps between steps.

Virtual Machine A virtual machine is a software environment that allows user code to have a consistent execution environment on many different computer platforms.

Definition: Agent-Based Modeling and Computer Languages

Agent-based modeling is a bottom-up approach to representing and investigating complex systems. Agent-based models can be implemented either computationally (e.g., through computer simulation) or non-computationally (e.g., with participatory simulation). The close match between the capabilities of available software and the requirements of agent-based modeling make these options a natural choice for many agent-based models. Of course, realizing the potential benefits of this natural match necessitates the use of computer languages to express the designs of agent-based models. A wide range of computer design and programming languages can play this role including both domain-specific and general-purpose languages. The domain-specific languages include business-oriented languages (e.g., spreadsheet programming tools), science and engineering languages (e.g., Mathematica), and dedicated agent-based modeling languages (e.g., NetLogo). The general-purpose languages can be used directly (e.g., Java programming) or within agent-based modeling toolkits (e.g., Repast). The choice that is most appropriate for each modeling project depends on both the requirements of that project and the resources available to implement it.

Agent-Based Modeling

This introduction follows Macal and North (2007). The term agent-based modeling (ABM) refers to the computational modeling of a system as comprised of a number of independent,

interacting entities, which are referred to as “agents.” Generally, an agent-based system is made up of agents that interact, adapt, and sustain themselves while interacting with other agents and adapting to a changing environment. The fundamental feature of an agent is its autonomy, the capability of the agent to act independently without the need for direction from external sources. Agent behaviors allow agents to take in information from their environment, which includes their interactions with other agents, process the information and make some decision about their next action, and take the action. Jennings (2000) provides a formal computer science-oriented view of agency emphasizing the essential characteristic of autonomous behavior.

Beyond the essential characteristic of autonomy, there is no universal agreement on the precise definition of the term “agent,” as used in agent-based modeling. Some consider any type of independent component, whether it be a software model or a software model of an extant individual, to be an agent (Bonabeau 2001). An independent component’s behaviors can be modeled as consisting of anything from simple reactive decision rules to multidimensional behavior complexes based on adaptive artificial intelligence (AI) techniques.

Other authors insist that a component’s behavior must be adaptive in order for the entity to be considered an agent. The agent label is reserved for components that can adapt to their environment, by learning from the successes and failures of their interactions with other agents, and change their behaviors in response. Casti (1997) argues that agents should contain both base-level rules for behavior and a higher-level set of “rules to change the rules.” The base-level rules provide responses to the environment while the “rules to change the rules” provide adaptation (Casti 1997).

From a practical modeling standpoint, agent characteristics can be summarized as follows:

- Agents are identifiable as self-contained individuals. An agent has a set of characteristics and rules governing its behaviors.
- Agents are autonomous and self-directed. An agent can function independently in its

environment and in its interactions with other agents, at least over a limited range of situations that are of interest.

- An agent is situated, living in an environment with which it interacts along with other agents. Agents have the ability to recognize and distinguish the traits of other agents. Agents also have protocols for interaction with other agents, such as for communication, and the capability to respond to the environment.
- An agent may be goal directed, having targets to achieve with respect to its behaviors. This allows an agent to compare the outcome of its behavior to its goals. An agent's goals need not be comprehensive or well defined. For example, an agent does not necessarily have formally stated objectives it is trying to maximize.
- An agent might have the ability to learn and adapt its behaviors based on its experiences. An agent might have rules that modify its behavior over time. Generally, learning and adaptation at the agent level requires some form of memory to be built into the agent's behaviors.

Often, in an agent-based model, the population of agents varies over time, as agents are born and die. Another form of adaptation can occur at the agent population level. Agents that are fit are better able to sustain themselves and possibly reproduce as time in the simulation progresses, while agents that have characteristics less suited to their continued survival are excluded from the population.

Another basic assumption of agent-based modeling is that agents have access only to local information. Agents obtain information about the rest of the world only through their interactions with the limited number of agents around them at any one time, and from their interactions with a local patch of the environment in which they are situated.

These aspects of how agent-based modeling treats agents highlight the fact that the full range of agent diversity can be incorporated into an agent-based model. Agents are diverse and heterogeneous as well as dynamic in their attributes and behavioral rules. There is no need to make agents homogeneous through aggregating agents

into groups or by identifying the “average” agent as representative of the entire population. Behavioral rules vary in their sophistication, how much information is considered in the agent decisions (i.e., cognitive “load”), the agent's internal models of the external world including the possible reactions or behaviors of other agents, and the extent of memory of past events the agent retains and uses in its decisions. Agents can also vary by the resources they have managed to accumulate during the simulation, which may be due to some advantage that results from specific attributes. The only limit on the number of agents in an agent-based model is imposed by the computational resources required to run the model.

As a point of clarification, agent-based modeling is also known by other names. ABS (agent-based systems), IBM (individual-based modeling), and MAS (multi-agent systems) are widely used acronyms, but “ABM” will be used throughout this discussion. The term “agent” has connotations other than how it is used in ABM. For example, ABM agents are different from the typical agents found in mobile agent systems. “Mobile agents” are lightweight software proxies that roam the World Wide Web and perform various functions programmed by their owners such as gathering information from Web sites. To this extent, mobile agents are autonomous and share this characteristic with agents in ABM.

Types of Computer Languages

A “computer language” is a method of specifying directives for computers. “Computer programming languages,” or “programming languages,” are an important category of computer languages. A programming language is a computer language that allows any computable activity to be expressed. This entry focuses on computer programming languages rather than the more general computer languages, since virtually all agent-based modeling systems require the power of programming languages. This entry sometimes uses the simpler term “computer languages” when referring to computer programming languages. According to Watson (1989),

Programming languages are used to describe algorithms, that is, sequences of steps that lead to the solution of problems. . . A programming language can be considered to be a ‘notation’ that can be used to specify algorithms with precision.

Watson (1989) goes on to say that “programming languages can be roughly divided into four groups: imperative languages, functional languages, logic programming languages, and others.” Watson (1989) states that in imperative languages “there is a fundamental underlying dependence on the assignment operation and on variables implemented as computer memory locations, whose contents can be read and altered.” However, “in functional languages (sometimes called applicative languages) the fundamental operation is function application” (Watson 1989). Watson cites LISP as an example. Watson (1989) continues by noting that “in a logic programming language, the programmer needs only to supply the problem specification in some formal form, as it is the responsibility of the language system to infer a method of solution.” Watson cites Prolog as an example.

A useful feature of most functional languages, many logic programming languages, and some imperative languages is higher-order programming. According to Reynolds (1998),

In analogy with mathematical logic, we will say that a programming language is higher-order if procedures or labels can occur as data, i.e., if these entities can be used as arguments to procedures, as results of functions, or as values of assignable variables. A language that is not higher-order will be called first-order.

Watson (1989) offers that “another way of grouping programming languages is to classify them as procedural or declarative languages.” Elaborating, Watson (1989) states that

Procedural languages . . . are those in which the action of the program is defined by a series of operations defined by the programmer. To solve a problem, the programmer has to specify a series of steps (or statements) which are executed in sequence

On the other hand, Watson (1989) notes that

Programming in a declarative language (or non-procedural language) involves the specification of a set of rules defining the solution to the problem; it is then up to the computer to determine

how to reach a solution consistent with the given rules. . . The language Prolog falls into this category, although it retains some procedural aspects. Another widespread non-procedural system is the spreadsheet program.

Imperative and functional languages are usually procedural, while logic programming languages are generally declarative. This distinction is important since it implies that most imperative and functional languages require users to define how each operation is to be completed, while logic programming languages only require users to define what is to be achieved. However, when faced with multiple possible solutions with different execution speeds and memory requirements, imperative and functional languages offer the potential for users to explicitly choose more efficient implementations over less efficient ones. Logic programming languages generally need to infer which solution is best from the problem description and may or may not choose the most efficient implementation. Naturally, this potential strength of imperative and functional languages may also be cast as a weakness. With imperative and functional language, users need to correctly choose a good implementation among any competing candidates that may be available.

Similarly to Watson (1989), Van Roy and Haridi (2004) define several common computational models, namely, those that are object oriented, those that are logic based, and those that are functional. Object-oriented languages are procedural languages that bind procedures (i.e., “encapsulated methods”) to their corresponding data (i.e., “fields”) in nested hierarchies (i.e., “inheritance” graphs) such that the resulting “classes” can be instantiated to produce executable instances (i.e., “objects”) that respond to multiple related messages (i.e., “polymorphism”). Logic-based languages correspond to Watson’s (1989) logic programming languages. Similarly, Van Roy and Haridi’s (2004) functional languages correspond to those of Watson (1989).

Two additional types of languages can be added to Van Roy and Haridi’s (2004) list of three. These are unstructured and structured languages (Dijkstra 1968). Both unstructured and structured languages are procedural languages.

Unstructured languages are languages that rely on step-by-step solutions such that the solutions can contain arbitrary jumps between steps (Dijkstra 1968). BASIC, COBOL, Fortran, and C are examples of unstructured languages. The arbitrary jumps are often implemented using “goto” statements. Unstructured languages were famously criticized by Edsger Dijkstra in his classic paper “Go To Statement Considered Harmful” (1968). This and related criticism led to the introduction of structured languages.

Structured languages are languages that divide programs into separate modules, each of which has one controlled entry point, a limited number of exit points, and no internal jumps (Dijkstra 1968). Following Stevens et al. (1974) “the term module is used to refer to a set of one or more contiguous program statements having a name by which other parts of the system can invoke it and preferably having its own distinct set of variable names.” Structured language modules, often called procedures, are generally intended to be small. As such, large numbers of them are usually required to solve complex problems. Standard Pascal is an example of structured, but not object-oriented, language. As stated earlier, C is technically an unstructured language (i.e., it allows jumps within procedures and “long jumps” between procedures), but it is used so often in a structured way that many people think of it as a structured language.

The quality of modularization in structured language code is often considered to be a function of coupling and cohesion (Stevens et al. 1974). Coupling is the tie between modules such that the proper functioning of one module depends on the functioning of another module. Cohesion refers to the ties within a module such that proper functioning of one line of code in a module depends on the functioning of another one line of code in the same module. The goal for modules is maximizing cohesion while minimizing coupling.

Object-oriented languages are a subset of structured languages. Object-oriented methods and classes are structured programming modules that have special features for binding data, inheritance, and polymorphism. The previously introduced concepts of coupling and cohesion apply to classes,

objects, methods, and fields the same way that they apply to generic structured language modules. Objective-C, C++, C#, and Java are all examples of object-oriented languages. As with C, the languages Objective-C, C++, and C# offer goto statements but they have object-oriented features and are generally used in a structured way. Java is an interesting case in that the word “goto” is reserved as a keyword in the language specification, but it is not intended to be implemented.

It is possible to develop agent-based models using any of the programming languages discussed above, namely, unstructured languages, structured languages, object-oriented languages, logic-based languages, and functional languages. Specific examples are provided later in this entry. However, certain features of programming languages are particularly well suited for supporting the requirements of agent-based modeling and simulation.

Requirements of Computer Languages for Agent-Based Modeling

The requirements of computer languages for agent-based modeling and simulation include the following:

- There is a need to create well-defined modules that correspond to agents. These modules should bind together agent state data and agent behaviors into integrated independently addressable constructs. Ideally these modules will be flexible enough to change structure over time and to optionally allow fuzzy boundaries to implement models that go beyond methodological individualism (Heath 2005).
- There is a need to create well-defined containers that correspond to agent environments. Ideally these containers will be recursively nestable or will otherwise support sophisticated definitions of containment.
- There is a need to create well-defined spatial relationships within agent environments. These relationships should include notions of abstract space (e.g., lattices), physical space (e.g., maps), and connectedness (e.g., networks).

- There is a need to easily set up model configurations such as the number of agents, the relationships between agents, the environmental details, and the results to be collected.
- There is a need to conveniently collect and analyze model results.

Each of the kinds of programming languages, namely, unstructured languages, structured languages, object-oriented languages, logic-based languages, and functional languages, can address these requirements.

Unstructured languages generally support procedure definitions which can be used to implement agent behaviors. They also sometimes support the collection of diverse data into independently addressable constructs in the form of data structures often called “records.” However, they generally lack support for binding procedures to individual data items or records of data items. This lack of support for creating integrated constructs also typically limits the language-level support for agent containers. Native support for implementing spatial environments is similarly limited by the inability to directly bind procedures to data.

As discussed in the previous section, unstructured languages offer statements to implement execution jumps. The use of jumps within and between procedures tends to reduce module cohesion and increase module coupling compared to structured code. The result is reduced code maintainability and extensibility compared to structured solutions. This is a substantial disadvantage of unstructured languages.

In contrast, many have argued that, at least theoretically, unstructured languages can achieve the highest execution speed and lowest memory usage of the language options since nearly everything is left to the application programmers. In practice, programmers implementing agent-based models in unstructured languages usually need to write their own tools to form agents by correlating data with the corresponding procedures. Ironically, these tools are often similar in design, implementation, and performance to some of the structured and object-oriented features discussed later.

Unstructured languages generally do not provide special support for application data

configuration, program output collection, or program results analysis. As such, these tasks usually need to be manually implemented by model developers.

In terms of agent-based modeling, structured languages are similar to unstructured languages in that they do not provide tools to directly integrate data and procedures into independently addressable constructs. Therefore, structured language support for agents, agent environments, and agent spatial relationships is similar to that provided by unstructured languages. However, the lack of jump statements in structured languages tends to increase program maintainability and extensibility compared to unstructured languages. This generally gives structured languages a substantial advantage over unstructured languages for implementing agent-based models.

Object-oriented languages build on the maintainability and extensibility advantages of structured languages by adding the ability to bind data to procedures. This binding in the form of classes provides a natural way to implement agents. In fact, object-oriented languages have their roots in Ole-Johan Dahl and Kristen Nygaard’s Simula simulation language (Dahl and Nygaard 1966, 2001; Van Roy and Haridi 2004)! According to Dahl and Nygaard (1966),

SIMULA (SIMULATION LAnguage) is a language designed to facilitate formal description of the layout and rules of operation of systems with discrete events (changes of state). The language is a true extension of ALGOL 60 (Backus et al. 1963), i.e., it contains ALGOL 60 as a subset. As a programming language, apart from simulation, SIMULA has extensive list processing facilities and introduces an extended co-routine concept in a high-level language.

Dahl and Nygaard go on to state the importance of specific languages for simulation (1966) as follows:

Simulation is now a widely used tool for analysis of a variety of phenomena: nerve networks, communication systems, traffic flow, production systems, administrative systems, social systems, etc. Because of the necessary list processing, complex data structures and program sequencing demands, simulation programs are comparatively difficult to write in machine language or in ALGOL or FORTRAN. This alone calls for the introduction of simulation languages.

However, still more important is the need for a set of basic concepts in terms of which it is possible to approach, understand and describe all the apparently very different phenomena listed above. A simulation language should be built around such a set of basic concepts and allow a formal description which may generate a computer program. The language should point out similarities and differences between systems and force the research worker to consider all relevant aspects of the systems. System descriptions should be easy to read and print and hence useful for communication.

Again, according to Dahl and Nygaard (2001),

SIMULA I (1962–1965) and Simula 67 (1967) are the two first object-oriented languages. Simula 67 introduced most of the key concepts of object-oriented programming: both objects and classes, subclasses (usually referred to as inheritance) and virtual procedures, combined with safe referencing and mechanisms for bringing into a program collections of program structures described under a common class heading (prefixed blocks).

The Simula languages were developed at the Norwegian Computing Center, Oslo, Norway by Ole-Johan Dahl and Kristen Nygaard. Nygaard's work in Operational Research in the 1950s and early 1960s created the need for precise tools for the description and simulation of complex man-machine systems. In 1961 the idea emerged for developing a language that both could be used for system description (for people) and for system prescription (as a computer program through a compiler). Such a language had to contain an algorithmic language, and Dahl's knowledge of compilers became essential. . . . When the inheritance mechanism was invented in 1967, Simula 67 was developed as a general programming language that also could be specialised for many domains, including system simulation.

Generally, object-oriented classes are used to define agent templates, and instantiated objects are used to implement specific agents. Agent environment templates and spatial relationship patterns are also typically implemented using classes. Recursive environment nesting and abstract spaces, physical spaces, and connectedness can all be represented in relatively straightforward ways. Instantiated objects are used to implement specific agent environments and spatial relationships in individual models. Within these models, model configurations are also commonly implemented as objects instantiated from one or more classes. However, as with unstructured and structured languages, object-oriented languages

generally do not provide special support for application data configuration, collection of outputs, or analysis of results. As such, these tasks usually need to be manually implemented by model developers. Regardless of this, the ability to bind data and procedures provides such a straightforward method for implementing agents that most agent-based models are written using object-oriented languages.

It should be noted that traditional object-oriented languages do not provide a means to modify class and object structures once a program begins to execute. Newer “dynamic” object-oriented languages such as Groovy (Koenig et al. 2007) offer this capability. This potentially allows agents to gain and lose attributes and methods during the execution of a model based on the flow of events in a simulation. This in turn offers the possibility of implementing modules with fuzzy boundaries that are flexible enough to change structure over time.

As discussed in the previous section, logic-based languages offer an alternative to the progression formed by unstructured, structured, and object-oriented languages. Logic-based languages can provide a form of direct support for binding data (e.g., asserted propositions) with actions (e.g., logical predicates), sometimes including the use of higher-order programming. In principle, each agent can be implemented as a complex predicate with multiple nested sub-terms. The sub-terms, which may contain unresolved variables, can then be activated and resolved as needed during model execution. Agent templates which are analogous to object-oriented classes can be implemented using the same approach but with a larger number of unresolved variables. Agent environments and the resulting relationships between agents can be formed in a similar way. Since each of these constructs can be modified at any time, the resulting system can change structure over time and may even allow fuzzy boundaries. In practice this approach is rarely, if ever, used. As with the previously discussed approaches, logic-based languages usually do not provide special support for application data configuration, output collection, or results analysis so these usually need to be manually developed.

Functional languages offer yet another alternative to the previously discussed languages. Like logic-based and object-oriented languages, functional languages often provide a form of direct support for binding data with behaviors. This support often leverages the fact that most functional languages support higher-order programming. As a result, the data is usually in the form of nested lists of values and functions, while the behaviors themselves are implemented in the form of functions. Agent templates (i.e., “classes”), agent environments, and agent relationships can be implemented similarly. Each of the lists can be dynamically changed during a simulation run so the model structure can evolve and can potentially have fuzzy boundaries. Unlike the other languages discussed so far, a major class of functional languages, namely, those designed for computational mathematics, usually include sophisticated support for program output collection and results analysis. An example is *Mathematica* (Wolfram 2013). If the application data is configured in mathematically regular ways, then these systems may also provide support for application data setup.

Example Computer Languages Useful for Agent-Based Modeling

Design Languages

Design languages provide a way to describe models at a more abstract level than typical programming languages. Some design languages ultimately offer the opportunity to compile to executable code. Other design languages act as intermediate stages between initial conceptualization and compliable implementation. In either case, the resulting design documents can be used to describe the model once they are complete.

Design Patterns

Patterns have offered a powerful yet simple way to conceptualize and communicate ideas in many disciplines since Christopher Alexander introduced them in the late 1970s (Alexander et al. 1977; Alexander 1979). Design patterns form a “common vocabulary” describing tried-and-true solutions for commonly faced software design problems

(Coplien 2001). Software design patterns were popularized by Gamma et al. (1995). They have subsequently been shown to be of substantial value in improving software quality and development efficiency. Several authors, such as North and Macal (2011), have suggested that there is great potential for patterns to improve the practice of agent-based modeling as well. North and Macal (2013) discussed product and process patterns. Product patterns are a vocabulary for designing or implementing models. Process patterns are methods for designing, implementing, or using models.

According to Alexander (1979), “each pattern is a three-part rule, which expresses a relation between a certain context, a problem, and a solution.” The first part of a pattern characterizes the situation in which the problem occurs. The second part defines the problem to be solved. The third part describes a resolution to the outstanding issue as well as its positive and negative consequences. Every pattern has both fixed and variable elements. The fixed elements define the pattern. The variable elements allow the pattern to be adapted for each situation. Each pattern identifies a set of decisions to make in the development of a system. Sets of patterns that have been adapted for the situation are then used as a vocabulary to describe solutions to problems. North and Macal (2013) introduce a catalog of patterns specifically for agent-based modeling. An example from North and Macal (2013) is shown in Table 1.

ODD Protocol

Grimm et al.’s (2006) ODD protocol describes models using a three-part approach: overview, concepts, and details. The model overview includes a statement of the model’s intent, a description of the main variables, and a discussion of the agent activities. The design concepts include a discussion of the foundations of the model. The details include the initial setup configuration, input value definitions, and descriptions of any embedded models. The resulting natural language document cannot be translated directly into executable code. However, it provides a basis for describing the design of models for publications, user documentation, and model development programmers.

Agent-Based Modeling and Computer Languages, Table 1 The scheduler scramble product design pattern

<i>Name:</i> Scheduler scramble
<i>Problem:</i> How can multiple agents act during the same scheduler pattern clock tick without biasing the model results or giving a long-term advantage or disadvantage to any one agent?
<i>Context:</i> Two or more agents from the agent-based model pattern may attempt to simultaneously execute behaviors during the same clock tick
<i>Forces:</i> Activating a behavior before other agents can be either an artificial advantage or disadvantage for the agent that goes first. Agent rules should not have to include coordination functions
<i>Solution:</i> The competing behaviors at each clock tick are scheduled in a random order. This is a simulation pattern
<i>Resulting context:</i> A sequential behavioral activation order with long-term fairness and that is unbiased is produced

UML

The Unified Modeling Language (UML) (Object Management Group 2001; Object Management Group 2013) is a predominantly visual approach to specifying the design of software. UML is particularly useful since it is flexible, general, and independent of particular programming languages. UML is widely supported by a wide range of software development environments (Object Management Group 2013). Some of these environments can compile fully specified UML diagrams to executable code.

The UML standard defines several types of diagrams that together provide a powerful and widely accepted approach to software design. These diagram types include use case diagrams, state diagrams, activity diagrams, class diagrams, and object diagrams (North and Macal 2007). A combination of these diagram types can be used to fully document both the underlying knowledge and the resulting designs of agent-based models.

Domain-Specific Languages

Domain-specific languages (DSLs) are computer languages that are highly customized to support a well-defined application area or “domain.” DSLs commonly include a substantial number of keywords that are nouns and verbs in the area of application as well as overall structures and execution patterns that correspond closely with the application area. DSLs are intended to allow users to write in a language that is closely aligned with their area of expertise.

DSLs often gain their focus by losing generality. For many DSLs, there are activities that can be

programmed in most computer languages that cannot be programmed in the given DSL. This is consciously done to simplify the DSL’s design and make it easier to learn and use. If a DSL is properly designed, then the loss of generality is often inconsequential for most uses since the excluded activities are chosen to be outside the normal range of application. However, even the best designed DSLs can occasionally be restrictive when the bounds of the language are encountered. Some DSLs provide special extension points that allow their users to program in a more general language such as C or Java when the limits of the DSL are reached. This feature is extremely useful, but requires more sophistication on the part of the user in that they need to know and simultaneously use both the DSL and the general language.

DSLs have the potential to implement specific features to support design patterns within a given domain. (North and Macal 2011, 2013) describe agent-based modeling design patterns in greater depth.

In principle, DSLs can be unstructured, structured, object oriented, logic based, or functional. In practice, DSLs are often structured languages or object-oriented languages and occasionally are functional languages. Commonly used ABM DSLs include business-oriented languages (e.g., spreadsheet programming tools), science and engineering languages (e.g., *Mathematica*), and dedicated agent-based modeling languages (e.g., NetLogo and Repast Symphony ReLogo).

Business Languages


Some of the most widely used business computer languages are those available in spreadsheet

packages. Spreadsheets are usually programmed using a “macro language.” As discussed further in North and Macal (2007), any modern spreadsheet program can be used to do basic agent-based modeling. The most common convention is to associate each row of a primary spreadsheet worksheet with an agent and use consecutive columns to store agent properties. Secondary worksheets are then used to represent the agent environment and to provide temporary storage for intermediate calculations. A simple loop is usually used to scan down the list of agents and to allow each one to execute in turn. The beginning and end of the scanning loop are generally used for special setup activities before and special cleanup activities after each round. An example agent spreadsheet from North and Macal (2007) is shown in Fig. 1. Agent spreadsheets have both strengths and weaknesses compared to the other ABM tools. Agent spreadsheets tend to be easy to build but they also tend to have limited capabilities. This balance makes spreadsheets ideal for


agent-based model exploration, scoping, and prototyping. Simple agent models can be implemented on the desktop using environments outside of spreadsheets as well (Fig. 2).

Science and Engineering Languages

Science and engineering languages embodied in commercial products such as *Mathematica*, *MATLAB*, *Maple*, and others can be used as a basis for developing agent-based models. Such systems usually have a large user base, are readily available on desktops, and are widely integrated into academic training programs. They can be used as rapid prototype development tools or as components of large-scale modeling systems. Science and engineering languages have been applied to agent-based modeling. Their advantages include a fully integrated development environment, their interpreted (as opposed to compiled) nature providing immediate feedback to users during the development process, and a packaged user interface. Integrated tools provide



Agent-Based Modeling and Computer Languages, Fig. 1 An example agent spreadsheet (North and Macal 2007)



Agent-Based Modeling and Computer Languages, Fig. 2 An example agent spreadsheet code (North and Macal 2007)

support for data import and graphical display. Macal (2004) describes the use of *Mathematica* and MATLAB in agent-based simulation, and Macal and Howe (2005) detail investigations into linking *Mathematica* and MATLAB to the Repast ABM toolkit to make use of Repast's simulation scheduling algorithms. In the following sections, we focus on MATLAB and *Mathematica* as representative examples of science and engineering languages.

MATLAB and *Mathematica* are both examples of Computational Mathematics Systems (CMS). CMS allow users to apply powerful mathematical algorithms to solve problems through a convenient and interactive user interface. CMS typically

supply a wide range of built-in functions and algorithms. MATLAB, *Mathematica*, and Maple are examples of commercially available CMS whose origins go back to the late 1980s. CMS are structured in two main parts: (1) the user interface that allows dynamic user interaction and (2) the underlying computational engine, or kernel, that performs the computations according to the user's instructions. Unlike conventional programming languages, CMS are interpreted instead of compiled, so there is immediate feedback to the user, but some performance penalty is paid. The underlying computational engine is written in the C programming language for these systems, but C coding is unseen by the user. The

most recent releases of CMS are fully integrated systems, combining capabilities for data input and export, graphical display, and the capability to link to external programs written in conventional languages such as C or Java using inter-process communication protocols. The powerful features of CMS, their convenience of use, the need to learn only a limited number of instructions on the part of the user, and the immediate feedback provided to users are features of CMS that make them good candidates for developing agent-based simulations.

A further distinction can be made among CMS. A subset of CMS are what is called Computational Algebra Systems (CAS). CAS are computational mathematics systems that calculate using symbolic expressions. CAS owe their origins to the LISP programming language, which was the earliest functional programming language (McCarthy 1960). Macsyma (www.scientek.com/macsyma) and Scheme (Springer and Freeman 1989) (www.swiss.ai.mit.edu/projects/scheme) are often mentioned as important implementations leading to present-day CAS. Typical uses of CAS are equation solving, symbolic integration and differentiation, exact calculations in linear algebra, simplification of mathematical expressions, and variable precision arithmetic. Computational mathematics systems consist of numeric processing systems or symbolic processing systems, or possibly a combination of both. Especially when algebraic and numeric capabilities are combined into a multi-paradigm programming environment, new modeling possibilities open up for developing sophisticated agent-based simulations with minimal coding.

Mathematica *Mathematica* is a commercially available numeric processing system with enormous integrated numerical processing capability (<http://www.wolfram.com>). Beyond numeric processing, *Mathematica* is a fully functional programming language. Unlike MATLAB, *Mathematica* is a symbolic processing system that uses term replacement as its primary operation. Symbolic processing means that variables can be used before they have values assigned to them; in contrast, a numeric processing language requires

that every variable have a value assigned to it before it is used in the program. In this respect, although *Mathematica* and MATLAB may appear similar and share many capabilities, *Mathematica* is fundamentally much different than MATLAB, with a much different style of programming and ultimately with a different set of capabilities applicable to agent-based modeling.


Mathematica's symbolic processing capabilities allow one to program in multiple programming styles, either as alternatives or in combination, such as functional programming, logic programming, procedural programming, and even object-oriented programming styles. Like MATLAB, *Mathematica* is also an interpreted language, with the kernel of *Mathematica* running in the background in C. In terms of data types, everything is an expression in *Mathematica*. An expression is a data type with a head and a list of arguments in which even the head of the expression is part of the expression's arguments.

The *Mathematica* user interface consists of a what is referred to as a notebook (Fig. 3). A *Mathematica* notebook is a fully integratable development environment and a complete publication environment. The *Mathematica* application programming interface (API) allows programs written in C, Fortran, or Java to interact with *Mathematica*. The API has facilities for dynamically calling routines from *Mathematica* as well as calling *Mathematica* as a computational engine.

Figure 3 shows *Mathematica* desktop notebook environment. A *Mathematica* notebook is displayed in its own window. Within a notebook, each item is contained in a cell. The notebook cell structure has underlying coding that is accessible to the user.

In *Mathematica*, a network representation consists of combining lists of lists, or more generally expressions of expressions, to various depths. For example, in *Mathematica*, an agent can be represented explicitly as an expression that includes a head named agent, a sequence of agent attributes, and a list of the agent's neighbors. Agent data and methods are linked together by the use of what are called "up values."

Example references for agent-based simulation using *Mathematica* include Gaylord and Davis (1999), Gaylord and Nishidate (1994), and




Agent-Based Modeling and Computer Languages, Fig. 3 Example *Mathematica* cellular automata model

Gaylord and Wellin (1995). Gaylord and D'Andria (1998) describe applications in social agent-based modeling.

MATLAB The MATrix LABoratory (MATLAB) is a numeric processing system with enormous integrated numerical processing capability (<http://www.mathworks.com>). It uses a scripting-language approach to programming. MATLAB is a high-level matrix/array language with control flow, functions, data structures, input/output, and object-oriented programming features. The user interface

consists of the MATLAB desktop, which is a fully integrated and mature development environment. MATLAB has an application programming interface (API). The MATLAB API allows programs written in C, Fortran, or Java to interact with MATLAB. There are facilities for calling routines from MATLAB (dynamic linking) as well as routines for calling MATLAB as a computational engine, as well as for reading and writing specialized MATLAB files.

Figure 4 shows the MATLAB desktop environment illustrating the Game of Life, which is a



Agent-Based Modeling and Computer Languages, Fig. 4 Example MATLAB cellular automata model

standard MATLAB demonstration. The desktop consists of four standard windows: a command window, which contains a command line, the primary way of interacting with MATLAB; the workspace, which indicates the values of all the variables currently existing in the session; a command history window that tracks the entered command; and the current directory window. Other windows allow text editing of programs and graphical output display.

When it comes to agent-based simulation, as in most types of coding, the most important indicator of the power of a language for modeling is the extent of and the sophistication of the allowed data types and data structures. As Sedgewick observes:

For many applications, the choice of the proper data structure is really the only major decision involved in the implementation; once the choice has been made only very simple algorithms are needed. (Sedgewick 1988)

The flexibility of data types plays an important role in developing large-scale, extensible models for agent-based simulation. In MATLAB the primary data type is the double array, which is

essentially a two-dimensional numeric matrix. Other data types include logical arrays, cell arrays, structures, and character arrays.

For agent-based simulations that define agent relationships based on networks, connectivity of the links defines the scope of agent interaction and locally available information. Extensions to modeling social networks require the use of more complex data structures than the matrix structure commonly used for grid representations. Extensions from grid topologies to network topologies are straightforward in MATLAB and similarly in *Mathematica*. In MATLAB, a network representation consists of combining cell arrays or structures in various ways.

The MATLAB desktop environment showing the Game of Life demonstration appears in Fig. 4. The Game of Life is a cellular automaton invented by mathematician John Conway that involves live and dead cells in cellular automata grid. In MATLAB, the agent environment is a sparse matrix that is initially set to all zeros. Whether cells stay alive, die, or generate new cells depends upon how many of their eight possible neighbors are alive. By using sparse matrices, the calculations

required become very simple. Pressing the “Start” button automatically seeds this universe with several small random communities and initiates a series of cell updates. After a short period of simulation, the initial random distribution of live (i.e., highlighted) cells develops into sets of sustainable patterns that endure for generations.

Several agent-based models using MATLAB have been published in addition to the Game of Life. These include a model of political institutions in modern Italy (Bhavnani 2003), a model of pair interactions and attitudes (Pearson and Boudarel 2001), a bargaining model to simulate negotiations between water users (Thoyer et al. 2001), and a model of sentiment and social mitosis based on Heider’s Balance Theory (Guetzkow et al. 1972; Wang and Thorngate 2003). The latter model uses Euler, a MATLAB-like language. Thorngate argues for the use of MATLAB as an important tool to teach simulation programming techniques (Thorngate 2000).

Dedicated Agent-Based Modeling Languages

Dedicated agent-based modeling languages are DSLs that are designed to specifically support agent-based modeling. Several such languages currently exist. These languages are functionally differentiated by the underlying assumptions their designers made about the structures of agent-based models. The designers of some of these languages assume quite a lot about the situations being modeled and use this information to provide users with pre-completed or template components. The designers of other languages make comparatively fewer assumptions and encourage users to implement a wider range of models. However, more work is often needed to build models in these systems. This entry will discuss two selected examples, namely, NetLogo and Repast Symphony flowcharts.

NetLogo

NetLogo is an education-focused ABM environment (Wilensky 1999). The NetLogo language uses a modified version of the Logo programming language (Harvey 1997). NetLogo itself is Java based and is free for use in education and research.


More information on NetLogo and downloads can be found at <http://ccl.northwestern.edu/netlogo/>.

NetLogo is designed to provide a basic computational laboratory for teaching complex adaptive systems concepts. NetLogo was originally developed to support teaching, but it can be used to develop a wider range of applications. NetLogo provides a graphical environment to create programs that control graphic “turtles” that reside in a world of “patches” that is monitored by an “observer.” NetLogo’s DSL is limited to its turtle and patch paradigm. However, NetLogo models can be extended using Java to provide for more general programming capabilities. An example NetLogo model of an ant colony (Wilensky 1999) (center) feeding on three food sources (upper left corner, lower left corner, and middle right) is shown in Fig. 5. Example code (Wilensky 1999) from this model is shown in Fig. 6.

Repast Symphony Flowcharts

The Recursive Porous Agent Simulation Toolkit (Repast) is a free and open source suit of agent-based modeling and simulation library (ROAD 2013). The Repast Suite is a family of advanced, free, and open source agent-based modeling and simulation software that have collectively been under continuous development for over 10 years. Repast Symphony is a richly interactive and easy to learn Java-based modeling environment that is designed for use on workstations and small computing clusters. Repast for high-performance computing (HPC) is a lean and expert-focused C++-based modeling library that is designed for use on large computing clusters and supercomputers. Repast Symphony and Repast HPC share a common architecture. Information on the Repast Suite and free downloads can be found at <http://repast.sourceforge.net/>.

Repast Symphony (North et al. 2013) includes advanced features for specifying, executing, and analyzing agent-based simulations. Repast Symphony offers several methods for specifying agents and agent environments including visual specification, specification with the dynamic object-oriented Groovy language (Koenig et al. 2007), and specification with Java. In principle, Repast Symphony’s visual DSL can be used for



Agent-Based Modeling and Computer Languages, Fig. 5 Example NetLogo ant colony model (Wilensky 1999)

any kind of programming, but models beyond a certain level of complexity are better implemented in Groovy or Java. As discussed later, Groovy and Java are general-purpose languages. All of Repast Symphony's languages can be fluidly combined in a single model. An example Repast Symphony zombie model is shown in Fig. 7 (North et al. 2013). In all cases, the user has a choice of a visual rich point-and-click interface or a "headless" batch interface to execute models (Fig. 8).


General Languages

Unlike DSLs, general languages are designed to take on any programming challenge. However, in order to meet this challenge, they are usually more complex than DSLs. This tends to make them more difficult to learn and use. Lahtinen et al. (2005)

document some of the challenges users face in learning general-purpose programming languages. Despite these issues, general-purpose programming languages are essential for allowing users to access the full capabilities of modern computers. Naturally, there are a huge number of general-purpose programming languages. This entry considers these options from two perspectives. First, general language toolkits are discussed. These toolkits provide libraries of functions to be used in a general-purpose host language. Second, the use of three raw general-purpose languages, namely, Java, C#, and C++, is discussed.

General Language Toolkits

As previously stated, general language toolkits are libraries that are intended be used in a general-



Agent-Based Modeling and Computer Languages, Fig. 6 Example NetLogo code from the ant colony model (Wilensky 1999)

purpose host language. These toolkits usually provide model developers with software for functions such as simulation time scheduling, results visualization, results logging, and model execution as well as domain-specific tools (North et al. 2006). Users of raw general-purpose languages have to write all of the needed features by themselves by hand.


A wide range of general language toolkits currently exist. This entry will discuss two selected examples, namely, Swarm and the Groovy and Java interfaces for Repast Simphony.

Swarm Swarm (Minar et al. 1996; Swarm Development Group 2013) is a free and open source agent-based modeling library. Swarm seeks to create a shared simulation platform for agent modeling and to facilitate the development of a wide range of models. Users build simulations by incorporating Swarm library components into their own programs. Information on Swarm


and free downloads can be found at <http://www.swarm.org/> From Marcus Daniels (Daniels 1999),

Swarm is a set of libraries that facilitate implementation of agent-based models. Swarm's inspiration comes from the field of Artificial Life. Artificial Life is an approach to studying biological systems that attempts to infer mechanism from biological phenomena, using the elaboration, refinement, and generalization of these mechanisms to identify unifying dynamical properties of biological systems...To help fill this need, Chris Langton initiated the Swarm project in 1994 at the Santa Fe Institute. The first version was available by 1996, and since then it has evolved to serve not only researchers in biology, but also anthropology, computer science, defense, ecology, economics, geography, industry, and political science.

The Swarm simulation system has two fundamental components. The core component runs general-purpose simulation code written in Objective-C, Tcl/Tk, and Java. This component handles most of the behind-the-scenes details. The external wrapper components run user-specific




Agent-Based Modeling and Computer Languages, Fig. 7 Example Repast Simphony zombie model (North et al. 2013)



Agent-Based Modeling and Computer Languages, Fig. 8 Example Repast Simphony visual behavior from a zombie model (North et al. 2013)

Agent-Based Modeling and Computer Languages,

Fig. 9 Example Swarm supply chain model (Swarm Development Group 2013)




simulation code written in either Objective-C or Java. These components handle most of the center stage work. An example Swarm supply chain model is shown in Fig. 9.

Repast Simphony Java and Groovy As previously discussed, Repast which is a suite is a free and open source of agent-based modeling and simulation software (ROAD 2013). Information on the Repast Suite and free downloads can be found at <http://repast.sourceforge.net/>. The Java-based Repast Simphony environment includes advanced features for specifying, executing, and analyzing agent-based simulations. An example Repast Simphony zombie model is shown in Fig. 7 (North et al. 2013).


Repast Simphony offers several intermixable methods for specifying agents and agent environments including visual specification, specification with the dynamic object-oriented Groovy language (Koenig et al. 2007), and specification with Java. The Groovy approach uses the dynamic object-oriented Groovy language as shown in Fig. 10. The Java approach for an example Zombie is shown in Fig. 11.

Java

Java (Foxwell 1999) is a widely used object-oriented programming language that was developed and is maintained by Oracle Corporation. Java is known for its widespread “cross-platform” availability on many different types of hardware and operating systems. This capability comes from Java’s use of a “virtual machine” that allows binary code or “bytecode” to have a consistent execution environment on many different computer platforms. A large number of tools are available for Java program development including the powerful Eclipse development environment (2013) and many supporting libraries. Java uses reflection and dynamic method invocation to implement a variant of higher-order programming. Reflection is used for runtime class structure examination, while dynamic method invocation is used to call newly referenced methods at runtime. Java’s object orientation, cross-platform availability, reflection, and dynamic method invocation along with newer features such as annotations for including metadata in compiled code, generics for generalizing class, and aspects to



Agent-Based Modeling and Computer Languages, Fig. 10 Example Repast Symphony Groovy code from a zombie model (North et al. 2013)



Agent-Based Modeling and Computer Languages, Fig. 11 Example Repast Symphony Java code from a zombie model North et al. 2013)

implement dispersed but recurrent tasks make it a good choice for agent-based model development.

C#

C# (Archer 2001) is an object-oriented programming language that was developed and is maintained by Microsoft. C# is one of many languages that can be used to generate Microsoft .NET Framework code or Common Intermediate Language (CIL). Like Java bytecode, CIL is run using a “virtual machine” that potentially gives it a consistent execution environment on different computer platforms. A growing number of tools are emerging to support C# development. C# and the Microsoft .NET Framework more generally are in principle cross-platform, but in practice they are mainly executed under Microsoft Windows.

The Microsoft .NET Framework provides for the compilation into CIL of many different languages such as C#, Managed C++, and Managed Visual Basic to name just a few. Once these languages are compiled to CIL, the resulting modules are fully interoperable. This allows users to conveniently develop integrated software using a mixture of different languages. Like Java, C# supports reflection and dynamic method invocation for higher-order programming. C#'s object orientation, multilingual integration, generics, attributes for including metadata in compiled code, aspects, reflection, and dynamic method invocation make it well suited for agent-based model development, particularly on the Microsoft Windows platform.

C++

C++ is a widely used object-oriented programming language that was created by Bjarne Stroustrup (Stroustrup 2008) at AT&T. C++ is widely noted for both its object-oriented structure and its ability to be easily compiled into native machine code. C++ gives users substantial access to the underlying computer but also requires substantial programming skills.

Most C++ compilers are actually more properly considered C/C++ compilers since they can compile non-object-oriented C code as well as object-oriented C++ code. This allows

sophisticated users the opportunity to highly optimize selected areas of model code. However, this also opens the possibility of introducing difficult-to-resolve errors and hard-to-maintain code. It is also more difficult to port C++ code from one computer architecture to another than it is for virtual machine-based languages such as Java.

C++ can use a combination of Runtime Type Identification (RTTI) and function pointers to implement higher-order programming. Similar to the Java approach, C++ RTTI can be used for runtime class structure examination, while function pointers can be used to call newly referenced methods at runtime. C++'s object orientation, RTTI, function pointers, and low-level machine access make it a reasonable choice for the development of extremely large or complicated agent-based models.

Future Directions

Future developments in computer languages could have enormous implications for the development of agent-based modeling. Some of the challenges of agent-based modeling for the future include (1) scaling up models to handle large numbers of agents running on distributed heterogeneous processors across the grid, (2) handling the large amounts of data generated by agent models and making sense out of it, and (3) developing user-friendly interfaces and modular components in a collaborative environment that can be used by domain experts with little or no knowledge of standard computer coding techniques. Visual and natural language development environments that can be used by nonprogrammers are continuing to advance but remain to be proven at reducing the programming burden. There are a variety of next steps for the development of computer languages for agent-based modeling including the further development of DSLs, increasing visual modeling capabilities, and the development of languages and language features that better support pattern-based development. DSLs are likely to become increasingly available as agent-based modeling grows into a wider range of domains. More agent-based modeling systems

are developing visual interfaces for specifying model structures and agent behaviors. Many of these visual environments are themselves DSLs. The continued success of agent-based modeling will likely yield an increasing number of design patterns. Supporting and even automating implementations of these patterns may form a natural source for new language features. Many of these new features are likely to be implemented within DSLs.

Bibliography

- Alexander C (1979) *The timeless way of building*. Oxford University Press, Oxford
- Alexander C, Ishikawa S, Silverstein M (1977) *A pattern language*. Oxford University Press, Oxford
- Archer T (2001) *Inside C#*. Microsoft Press, Redmond
- Backus J, Bauer F, Green J, Katz C, McCarthy J, Naur P, Perlis A, Rutishauser H, Samuelson K, Vauquois B, Wegstein J, van Wijngaarden A, Woodger M (1963) Revised report on the algorithmic language ALGOL 60. In: Naur P (ed) *Communications of the association for computing machinery (ACM)*, vol 6. ACM, New York, pp 1–17
- Bhavnani R (2003) Adaptive agents, political institutions and civic traditions in modern Italy. *JASSS* 6(4). Available at <http://jasss.soc.surrey.ac.uk/6/4/1.html>
- Bonabeau E (2001) Agent-based modeling: methods and techniques for simulating human systems. *Proc Natl Acad Sci* 99(3):7280–7287
- Casti J (1997) *Would-be worlds: how simulation is changing the world of science*. Wiley, New York
- Coplien J (2001) Software patterns home page. Available at <http://hillside.net/patterns/>
- Dahl O-J, Nygaard K (1966) SIMULA – an ALGOL-based simulation language. *Commun ACM* 9:671–678
- Dahl O-J, Nygaard K (2001) How object-oriented programming started. Available at http://heim.ifi.uio.no/~kristen/FORSKNINGSOK_MAPPE/F_OO_start.html
- Daniels M (1999) Integrating simulation technologies with swarm. In: *Proceedings of the agent 1999 workshop on agent simulation: applications, models, and tools*. Argonne National Laboratory, Argonne
- Dijkstra E (1968) Go to statement considered harmful. *Commun ACM* 11(3):147–148
- Eclipse (2013) Eclipse home page. Available at <http://www.eclipse.org/>
- Foxwell H (1999) *Java 2 software development kit*. Linux J
- Gamma E, Helm R, Johnson R, Vlissides J (1995) *Design patterns: elements of reusable object-oriented software*. Addison-Wesley, Wokingham
- Gaylord R, D’Andria L (1998) *Simulating society: a Mathematica toolkit for modeling socioeconomic behavior*. Springer/TELOS, New York
- Gaylord R, Davis J (1999) Modeling nonspatial social interactions. *Math Educ Res* 8(2):1–4
- Gaylord R, Nishidate K (1994) *Modeling nature: cellular automata simulations with Mathematica*. Springer, New York
- Gaylord R, Wellin P (1995) *Computer simulations with Mathematica: explorations in complex physical and biological systems*. Springer/TELOS, New York
- Grimm V et al (2006) A standard protocol for describing individual-based and agent-based models. *Ecol Model* 198(1–2):115–126
- Guetzkow H, Kotler P, Schultz R (eds) (1972) *Simulation in social and administrative science*. Prentice Hall, Englewood Cliffs
- Harvey B (1997) *Computer science logo style*. MIT Press, Boston
- Heath J (2005) Methodological individualism. In: Zalta E (ed) *Stanford encyclopedia of philosophy*. Stanford University, Stanford. Available at <http://plato.stanford.edu/>
- Jennings N (2000) On agent-based software engineering. *Artif Intel* 117:277–296
- Koenig D, Glover A, King P, Laforge G, Skeet J (2007) *Groovy in action*. Manning Publications, Greenwich
- Lahtinen E, Ala-Mutka K, Jarvinen H-M (2005) A study of the difficulties of novice programmers. In: *Proceedings of the 10th annual SIGCSE conference on innovation and technology in computer science education*. ACM, Caparica
- Macal C (2004) Agent-based modeling and social simulation with mathematica and MATLAB. In: Macal C, Sallach D, North M (eds) *Proceedings of the agent 2004 conference on social dynamics: interaction, reflexivity and emergence*. Argonne National Laboratory, Argonne
- Macal C, Howe T (2005) Linking repast to computational mathematics systems: Mathematica and MATLAB. In: Macal C, Sallach D, North M (eds) *Proceedings of the agent 2005 conference on generative social processes, models, and mechanisms*. Argonne National Laboratory, Argonne
- Macal C, North M (2007) Agent-based modeling and simulation: desktop ABMS. In: Henderson SG, Biller B, Hsieh M-H, Shortle J, Tew JD, Barton RR (eds) *Proceedings of the 2007 winter simulation conference*. IEEE/ACM, Washington, DC
- McCarthy J (1960) Recursive functions of symbolic expressions and their computation by machine I. *J ACM* 3:184–195, ACM, New York, NY, USA
- Minar N, Burkhart R, Langton C, Askenazi M (1996) The swarm simulation system: a toolkit for building multi-agent simulations. Available at <http://alumni.media.mit.edu/~nelson/research/swarm/>
- North M, Macal C (2007) *Managing business complexity: discovering strategic solutions with agent-based modeling and simulation*. Oxford University Press: New York, NY, USA
- North M, Macal C (2011) Product design patterns for agent-based modeling. In: Jain S, Creasey R, Himmelsbach J (eds) *Proceedings of the 2011 winter simulation conference*. IEEE/ACM, Phoenix

- North M, Macal C (2013) Product and process patterns for agent-based modeling and simulation. *J Simulat* 8:25–36
- North M, Collier N, Vos R (2006) Experiences creating three implementations of the repast agent modeling toolkit. In: *ACM transactions on modeling and computer simulation*, vol 16, Issue 1. ACM, New York, pp 1–25
- North M, Collier N, Ozik J, Tatara E, Altaweel M, Macal M, Bragen M, Sydelko P (2013) Complex adaptive systems modeling with repast symphony. In: *Complex adaptive systems modeling*. Springer, Heidelberg
- Object Management Group (2001) OMG unified modeling language specification version 1.5. Object Management Group, Needham
- Object Management Group (2013) Object management group UML home page. Object Management Group, Needham
- Pearson D, Boudarel M-R (2001) Pair interactions: real and perceived attitudes. *JASSS* 4(4). Available at <http://www.soc.surrey.ac.uk/JASSS/4/4/4.html>
- Reynolds J (1998) Definitional interpreters for higher-order programming. In: *Higher-order and symbolic computation*. Kluwer, Dordrecht, pp 363–397
- ROAD (2013) Repast home page. Available at <http://repast.sourceforge.net/>
- Sedgewick R (1988) *Algorithms*, 2nd edn. Addison-Wesley, Reading, p 657
- Springer G, Freeman D (1989) *Scheme and the art of programming*. McGraw-Hill, New York
- Stevens W, Meyers G, Constantine L (1974) Structured design. *IBM Syst J* 13(2):115
- Stroustrup B (2008) Bjarne Stroustrup's FAQ. Available at http://www.research.att.com/~bs/bs_faq.html#invention
- Swarm Development Group home page (2013) Available at <http://www.swarm.org/>
- Thorngate W (2000) Teaching social simulation with MATLAB. *JASSS* 3(1). Available at <http://www.soc.surrey.ac.uk/JASSS/3/1/forum/1.html>
- Thoyer S, Morardet S, Rio P, Simon L, Goodhue R, Rausser G (2001) A bargaining model to simulate negotiations between water users. *JASSS* 4(2). Available at <http://www.soc.surrey.ac.uk/JASSS/4/2/6.html>
- Van Roy P, Haridi S (2004) *Concepts, techniques, and models of computer programming*. MIT Press, Cambridge
- Wang Z, Thorngate W (2003) Sentiment and social mitosis: implications of Heider's balance theory. *JASSS* 6(3). Available at <http://jasss.soc.surrey.ac.uk/6/3/2.html>
- Watson D (1989) *High-level languages and their compilers*. Addison-Wesley, Wokingham
- Wilensky U (1999) NetLogo <http://ccl.northwestern.edu/netlogo/>. Center for connected learning and computer-based modeling, Northwestern University, Evanston
- Wolfram Research (2013) Mathematica home page. Available at <http://www.wolfram.com/>

Computer Graphics and Games, Agent-Based Modeling in

Brian Mac Namee
School of Computing, Dublin Institute of
Technology, Dublin, Ireland

Article Outline

Glossary
Definition of the Subject
Introduction
Agent-Based Modelling in Computer Graphics
Agent-Based Modelling in CGI for Movies
Agent-Based Modelling in Games
Future Directions
Bibliography

Glossary

Computer generated imagery (CGI) The use of computer generated images for special effects purposes in film production.

Intelligent agent A hardware or (more usually) software-based computer system that enjoys the properties autonomy, social ability, reactivity and pro-activeness.

Non-player character (NPC) A computer controlled character in a computer game – as opposed to a player controlled character.

Virtual character A computer generated character that populates a virtual world.

Virtual world A computer generated world in which places, objects and people are represented as graphical (typically three dimensional) models.

Definition of the Subject

As the graphics technology used to create virtual worlds has improved in recent years, more and

more importance has been placed on the behavior of virtual characters in applications such as games, movies and simulations set in these virtual worlds simulations. The behavior of these virtual characters should be believable in order to create the illusion that virtual worlds are populated with living characters. This has led to the application of agent-based modeling to the control of virtual characters. There are a number of advantages of using agent-based modeling techniques which include the fact that they remove the requirement for hand controlling all agents in a virtual environment, and allow agents in games to respond to unexpected actions by players or users.

Introduction

Advances in computer graphics technology in recent years have allowed the creation of realistic and believable virtual worlds. However, as such virtual worlds have been developed for applications spanning games, education and movies it has become apparent that in order to achieve real believability, virtual worlds must be populated with life-like virtual characters. This is where the application of agent-based modeling has found a niche in the areas of computer graphics and, in a huge way, computer games. Agent-based modeling is a perfect solution to the problem of controlling the behaviors of the virtual characters that populate a virtual world. In fact, because virtual characters are embodied and autonomous these applications require an even stronger notion of agency than many other areas in which agent-based modeling is employed.

Before proceeding any further, and because there are so many competing alternatives, it is worth explicitly stating the definition of an intelligent agent that will inform the remainder of this article. Taken from Wooldridge and Jennings (1995) an intelligent agent is defined as “... *a hardware or (more usually) software-based computer system that enjoys the following properties:*

- **autonomy:** *agents operate without the direct intervention of humans or others, and have some kind of control over their actions and internal state;*
- **social ability:** *agents interact with other agents (and possibly humans) via some kind of agent-communication language;*
- **reactivity:** *agents perceive their environment, (which may be the physical world, a user via a graphical user interface, a collection of other agents, the INTERNET, or perhaps all of these combined), and respond in a timely fashion to changes that occur in it;*
- **pro-activeness:** *agents do not simply act in response to their environment, they are able to exhibit goal-directed behavior by taking the initiative."*

Virtual characters implemented using agent-based modeling techniques satisfy all of these properties. The characters that populate virtual worlds should be fully autonomous and drive their own behaviors (albeit sometimes following the orders of a director or player). Virtual characters should be able to interact believably with other characters and human participants. This property is particularly strong in the case of virtual characters used in games which by their nature are particularly interactive. It is also imperative that virtual characters appear to perceive their environments and react to events that occur in that environment, especially the actions of other characters or human participants. Finally virtual characters should be pro-active in their behaviors and not always require prompting from a human participant in order to take action.

The remainder of this article will proceed as follows. Firstly, a broad overview of the use of agent-based modeling in computer graphics will be given, focusing in particular on the genesis of the field. Following on from this, the focus will switch to the use of agent-based modeling techniques in two particular application areas: computer generated imagery (CGI) for movies, and computer games. CGI has been used to astounding effect in movies for decades, and in recent times has become heavily reliant on agent-based

modeling techniques in order to generate CGI scenes containing large numbers of computer generated extras. Computer games developers have also been using agent-based modeling techniques effectively for some time now for the control of non-player characters (NPCs) in games. There is a particularly fine match between the requirements of computer games and agent-based modeling due to the high levels of interactivity required.

Finally, the article will conclude with some suggestions for the future directions in which agent-based modeling technology in computer graphics and games is expected to move.

Agent-Based Modelling in Computer Graphics

The serious use of agent-based modeling in computer graphics first arose in the creation of autonomous groups and crowds - for example, crowds of people in a town square or hotel foyer, or flocks of birds in an outdoor scene. While initially this work was driven by visually unappealing simulation applications such as fire safety testing for buildings (Takahashi 1992), focus soon turned to the creation of visually realistic and believable crowds for applications such as movies, games and architectural walkthroughs. Computer graphics researchers realized that creating scenes featuring large virtual crowds by hand (a task that was becoming important for the applications already mentioned) was laborious and time-consuming and that agent-based modeling techniques could remove some of the animator's burden. Rather than requiring that animators hand-craft all of the movements of a crowd, agent-based systems could be created in which each character in a crowd (or flock, or swarm) could drive its own behavior. In this way the behavior of a crowd would emerge from the individual actions of the members of that crowd.

Two of the earliest, and seminal, examples of such systems are Craig Reynolds' Boids system (Reynolds 1987) and Tu & Terzopoulos' animations of virtual fish (Terzopoulos et al. 1994). The Boids system simulates the flocking behaviors exhibited in nature by schools of fish, or flocks

of birds. The system was first presented at the prestigious SIGGRAPH conference (www.siggraph.org) in 1987 and was accompanied by the short movie “Stanley and Stella in: Breaking the Ice”. Taking influence from the area of artificial life (or aLife) (Thalman and Thalman 1994), Reynolds postulated that the individual members of a flock would not be capable of complex reasoning, and so flocking behavior must emerge from simple decisions made by individual flock members. This notion of emergent behavior is one of the key characteristics of aLife systems.

In the original Boids system, each virtual agent (represented as a simple particle and known as a *boid*) used just three rules to control its movement. These were separation, alignment and cohesion and are illustrated in Fig. 1. Based on just these three simple rules extremely realistic flocking behaviors emerged. This freed animators from the laborious task of hand-scripting the behavior of each creature within the flock and perfectly demonstrates the advantage offered by agent-based modeling techniques for this kind of application.


The system created by Tu and Terzopoulos took a more complex approach in that they created complex models of biological fish. Their models took into account fish physiology, with a complex model of fish muscular structure, along with a perceptual model of fish vision. Using these they created sophisticated simulations in which properties such as schooling and predator avoidance were displayed. The advantage of this approach was that it was possible to create unique, unscripted, realistic simulations without the

intervention of human animators. Terzopoulos has since gone on to apply similar techniques to the control of virtual humans (Shao and Terzopoulos 2005).

Moving from animals to crowds of virtual humans, the Virtual Reality Lab at the Ecole Polytechnique Fédérale de Lausanne in Switzerland (vrlab.epfl.ch) led by Daniel Thalman has been at the forefront of this work for many years. They group currently has a highly evolved system, ViCrowd, for the animation of virtual crowds (Musse and Thalman 2001) which they model as a hierarchy which moves from individuals to groups to crowds. This hierarchy is used to avoid some of the complications which arise from trying to model large crowds in real time – one of the key goals of ViCrowd.

Each of the levels in the ViCrowd hierarchy can be modeled as an agent and this is done based on *beliefs*, *desires* and *intentions*. The beliefs of an agent represent the information that the agent possesses about the world, including information about places, objects and other agents. An agent’s desires represent the motivations of the agent regarding objectives it would like to achieve. Finally, the intentions of an agent represent the actions that an agent has chosen to pursue. The belief-desire-intention (BDI) model of agency was proposed by Rao and Georgeff (1991) and has been used in many other application areas of agent-based modeling.

ViCrowd has been used in ambitious applications including the simulation of a virtual city comprised of, amongst other things, a train station a park and a theater (Farenc et al. 2000). In all of



Computer Graphics and Games, Agent-Based Modeling in, Fig. 1 The three rules used by Reynolds’ original Boids system to simulate flocking behaviors

these environments the system was capable of driving the believable behaviors of large groups of characters in real-time.

It should be apparent to readers from the examples given thus far that the use of agent-based modeling techniques to control virtual characters gives rise to a range of unique requirements when compared to the use of agent based modeling in other application areas. The key to understanding these is to realize that the goal in designing agents for the control of virtual characters is typically not to design the most efficient or effective agent, but rather to design the most interesting or believable character. Outside of very practical applications such as evacuation simulations, when creating virtual characters, designers are concerned with maintaining what Disney, experts in this field, refer to as the *illusion of life* (Johnston and Thomas 1995).

This refers to the fact that the user of a system must believe that virtual characters are living, breathing creatures with goals, beliefs, desires, and, essentially, lives of their own. Thus, it is not so important for a virtual human to always choose the most efficient or cost effective option available to it, but rather to always choose reasonable actions and respond realistically to the success or failure of these actions. With this in mind, and following a similar discussion given in Isbister and Doyle (2002), some of the foremost researchers in virtual character research have the following to say about the requirements of agents as virtual characters.

Loyall writes (Loyall 1997) that “*Believable agents are personality-rich autonomous agents with the powerful properties of characters from the arts.*” Coming from a dramatic background it is not surprising that Loyall’s requirements reflect this. Agents should have strong personality and be capable of showing emotion and engaging in meaningful social relationships.

According to Blumberg (1996), “... *an autonomous animated creature is an animated object capable of goal-directed and time-varying behavior*”. The work of Blumberg and his group is very much concerned with virtual creatures, rather than humans in particular, and his requirements reflect this. Creatures must appear to make choices which

improve their situation and display sophisticated and individualistic movements.

Hayes–Roth and Doyle focus on the differences between “*animate characters*” and traditional agents (Hayes-Roth and Doyle 1998). With this in mind they indicate that agents’ behaviors must be “*variable rather than reliable*”, “*idiosyncratic instead of predictable*”, “*appropriate rather than correct*”, “*effective instead of complete*”, “*interesting rather than efficient*”, and “*distinctively individual as opposed to optimal*”.

Perlin and Goldberg (1996) concern themselves with building believable characters “*that respond to users and to each other in real-time, with consistent personalities, properly changing moods and without mechanical repetition, while always maintaining an author’s goals and intentions*”.

Finally, in characterizing believable agents, Bates (1992a) is quite forgiving requiring “*only that they not be clearly stupid or unreal*”. Such broad, shallow agents must “*exhibit some signs of internal goals, reactivity, emotion, natural language ability, and knowledge of agents ... as well as of the ... micro-world*”.

Considering these definitions, Isbister and Doyle (2002) identify the fact that the consistent themes which run through all of the requirements given above match the general goals of agency – virtual humans must display autonomy, reactivity, goal driven behavior and social ability – and again support the use of agent-based modeling to drive the behavior of virtual characters.

The Spectrum of Agents

The differences between the systems mentioned in the previous discussion are captured particularly well on the *spectrum of agents* presented by Aylett and Luck (2000). This positions agent systems on a spectrum based on their capabilities, and serves as a useful tool in differentiating between the various systems available. One end of this spectrum focuses on *physical agents* which are mainly concerned with simulation of believable physical behavior, (including sophisticated physiological models of muscle and skeleton systems), and of sensory systems. Interesting work at this end of the spectrum includes Terzopoulos’ highly

realistic simulation of fish (Terzopoulos et al. 1994) and his virtual stuntman project (Faloutsos et al. 2001) which creates virtual actors capable of realistically synthesizing a broad repertoire of lifelike motor skills.

Cognitive agents inhabit the other end of the agent spectrum and are mainly concerned with issues such as reasoning, decision making, planning and learning. Systems at this end of the spectrum include Funge's cognitive modeling approach (Funge 1999) which uses the situation calculus to control the behavior of virtual characters, and Nareyek's work on planning agents for simulation (Nareyek 2001), both of which will be described later in this article.

While the systems mentioned so far sit comfortably at either end of the agent spectrum, many of the most effective inhabit the middle ground. Amongst these are *c4* (Burke et al. 2002), used to great effect to simulate a virtual sheep dog with the ability to learn new behaviors, *Improv* (Perlin and Goldberg 1996) which augments sophisticated physical human animation with scripted behaviors and the *ViCrowd* system (Musse and Thalmann 2001) which sits on top of a realistic virtual human animation system and uses planning to control agents' behavior.

Virtual Fidelity

The fact that so many different agent-based modeling systems, for the control of virtual humans exist gives rise to the question why? The answer to this lies in the notion of *virtual fidelity*, as described by Badler (Badler et al. 1999). Virtual fidelity refers to the fact that virtual reality systems need only remain true to actual reality in so much as this is required by, and improves, the system.

In Määta (2002) the point is illustrated extremely effectively. The article explains that when game designers are architecting the environments in which games are set, the scale to which these environments are created is not kept true to reality. Rather, to ease players' movement in these worlds, areas are designed to a much larger scale, compared to character sizes, than in the real world. However, game players do not notice this digression from reality, and in fact

have a negative response to environments that are designed to be more true to life finding them cramped. This is a perfect example of how, although designers stay true to reality for many aspects of environment design, the particular blend of virtual fidelity required by an application can dictate certain real world restrictions can be ignored in virtual worlds.

With regard to virtual characters, virtual fidelity dictates that the set of capabilities which these characters should display is determined by the application which they are to inhabit. So, the requirements of an agent-based modeling system for CGI in movies would be very different to those of a agent-based modeling system for controlling the behaviors of game characters.

Agent-Based Modelling in CGI for Movies

With the success of agent-based modeling techniques in graphics firmly established there was something of a search for application areas to which they could be applied. Fortunately, the success of agent-based modeling techniques in computer graphics was paralleled with an increase in the use of CGI in the movie industry, which offered the perfect opportunity. In many cases CGI techniques were being used to replace traditional methods for creating expensive, or difficult to film scenes. In particular, scenes involving large numbers of people or animals were deemed no longer financially viable when set in the real world. Creating these scenes using CGI involved painstaking hand animation of each character within a scene, which again was not financially viable.

The solution that agent-based modeling offers is to make each character within a scene an intelligent agent that drives its own behavior. In this way, as long as the initial situation is set up correctly scenes will play out without the intervention of animators. The facts that animating for movies does not need to be performed in real-time, and is in no way interactive (there are no human users involved in the scene), make the use

of agent-based modeling a particularly fine match for this application area.

Craig Reynolds' Boids system (Reynolds 1987) which simulates the flocking behaviors exhibited in nature by schools of fish, or flocks of birds and was discussed previously is one of the seminal examples of agent-based modeling techniques being used in movie CGI. Reynold's approach was first used for CGI in the 1999 film "Batman Returns" (Burton 1992) to simulate colonies of bats. Reynold's technologies have been used in "The Lion King" (Allers and Minkoff 1994) and "From Dusk 'Till Dawn" (Rodriguez 1996) amongst other films. Reynolds' approach was so successful, in fact, that he was awarded an Academy Award for his work in 1998.

Similar techniques to those utilized in the Boids system have been used in many other films to animate such diverse characters as ants, people and stampeding wildebeest. Two productions which were released in the same year, "Antz" (Darnell and Johnson 1998) by Dreamworks and "A Bug's Life" (Lasseter and Stanton 1998) by Pixar took great steps in using CGI effects to animate large crowds for. For "Antz" systems were developed which allowed animators easily create scenes containing large numbers of virtual characters modeling each as an intelligent agent capable of obstacle avoidance, flocking and other behaviors. Similarly, the creators of "A Bug's Life" created tools which allowed animators easily combine pre-defined motions (known as alibis) to create behaviors which could easily be applied to individual agents in scenes composed of hundreds of virtual characters.

However, the largest jump in the use of agent-based modeling in movie CGI was made in the recent Lord of the Rings trilogy (Jackson 2001, 2002, 2003). In these films the bar was raised markedly in terms of the sophistication of the virtual characters displayed and the sheer number of characters populating each scene. To achieve the special effects shots required by the makers of these films, the Massive software system was developed by Massive Software (www.massivesoftware.com). This system (Aitken et al. 2004; Koeppel 2002) uses agent-based modeling techniques, again inspired by aLife, to create

virtual extras that control their own behaviors. This system was put to particularly good use in the large scale battle sequences that feature in all three of the Lord of the Rings films. Some of the sequences in the final film of the trilogy, the Return of the King, contain over 200,000 digital characters.

In order to create a large battle scene using the Massive software, each virtual extra is represented as an intelligent agent, making its own decisions about which actions it will perform based on its perceptions of the world around it. Agent control is achieved through the use of fuzzy logic based controllers in which the state of an agent's brain is represented as a series of motivations, and knowledge it has about the world – such as the state of the terrain it finds itself on, what kinds of other agents are around it and what these other agents are doing. This knowledge about the world is perceived through simple simulated visual, auditory and tactile senses. Based on the information they perceive agents decide on a best course of action. Designing the brains of these agents is made easier that it might seem at first by the fact that agents are developed for short sequences, and so a small range of possible tasks. So for example, separate agent models would be used for a fighting scene and a celebration scene.

In order to create a large crowd scene using Massive animators initially set up an environment populating it with an appropriate cast of virtual characters where the brains of each character are slight variations (based on physical and personality attributes) of a small number of archetypes. The scene will then play itself out with each character making it's own decisions. Therefore there is no need for any hand animation of virtual characters. However, directors can view the created scenes and by tweaking the parameters of the brains of the virtual characters have a scene play out in the exact way that they require.

Since being used to such impressive effect in the Lord of the Rings trilogy (the developers of the Massive system were awarded an academy award for their work), the Massive software system has been used in numerous other films such as "I, Robot" (Proyas 2004), "The Chronicles of Narnia: The Lion, the Witch and the Wardrobe"

(Adamson 2005) and “Ratatouille” (Bird and Pinkava 2007) along with numerous television commercials and music videos.

While the achievements of using agent-based modeling for movie CGI are extremely impressive, it is worth noting that none of these systems run in real-time. Rather, scenes are rendered by banks of high powered computers, a process that can take hours for relatively simple scenes. For example, the famous Prologue battle sequence in the “Lord of the Rings: The Fellowship of the Ring” took a week to render. When agent-based modeling is applied to the real-time world of computer games, things are very different.

Agent-Based Modelling in Games

Even more so than in movies, agent-based modeling techniques have been used to drive the behaviors of virtual characters in computer games. As games have become graphically more realistic (and in recent years they have become extremely so) game-players have come to expect that games are set in hugely realistic and believable virtual worlds. This is particularly evident in the widespread use of realistic physics modeling which is now commonplace in games (Sánchez-Crespo 2006). In games that make strong use of physics modeling, objects in the game world topple over when pushed, float realistically when dropped in water and generally respond as one would expect them to. Players expect the same to be true of the virtual characters that populate virtual game worlds. This can be best achieved by modeling virtual characters as embodied virtual agents. However, there are a number of constraints which have a major influence on the use of agent-based modeling techniques in games.

The first of these constraints stems from the fact that modern games are so highly interactive. Players expect to be able to interact with all of the characters they encounter within a game world. These interactions can be as simple as having something to shoot at or having someone to race against; or involve much more sophisticated interactions in which a player is expected to converse with a virtual character to find out specific

information or to cooperate with a virtual character in order to accomplish some task that is key to the plot of a game. Interactivity raises a massive challenge for practitioners as there is very little restriction in terms of what the player might do. Virtual characters should respond in a believable way at all times regardless of how bizarre and unexpected the actions of the player might be.

The second challenge comes from the fact that the vast majority of video games should run in real time. This means that the computational complexity must be kept to a reasonable level as there are only a finite number of processor cycles available for AI processing. This problem is magnified by the fact that an enormous amount of CPU power is usually dedicated to graphics processing. When compared to the techniques that can be used for controlling virtual characters in films some of the techniques used in games are rudimentary due to this real-time constraint.

Finally, modern games resemble films in the fact that creators go to great lengths to include intricate storylines and control the building of tension in much the way that film script writers do. This means that games are tested heavily in order to ensure that the game proceeds smoothly and that the level of difficulty is finely tuned so as to always hold the interest of a player. In fact, this testing of games has become something of a science in itself (Thompson 2007). Using autonomous agents gives game characters the ability to do things that are unexpected by the game designers and so upset their well laid plans. This can often be a barrier to the use of sophisticated techniques such as learning. Unfortunately there is also a barrier to the discussion of agent-based modeling techniques used in commercial games. Because of the very competitive nature of the games industry, game development houses often consider the details of how their games work as valuable trade secrets to be kept well guarded. This can make it difficult to uncover the details of how particularly interesting features of a game are implemented. While this situation is improving – more commercial game developers are speaking at games conferences about how their games are developed and the release of game systems development kits for the development of game modifications (or *mods*) allows researchers to plumb

the depths of game code – it is still often impossible to find out the implementation details of very new games.

Game Genres

Before discussing the use of agent-based modeling in games any further, it is worth making a short clarification on the kinds of computer games that this article refers to. When discussing modern computer games, or video games, this article does not refer to computer implementations of traditional games such as chess, backgammon or card games such as solitaire. Although these games are of considerable research interest (chess in particular has been the subject of extremely successful research (Feng-Hsiung 2002)) they are typically not approached using agent-based modeling techniques. Typically, artificial intelligence approaches to games such as these rely largely on sophisticated searching techniques which allow the computer player to search through a multitude of possible future situations dictated by the moves it will make and the moves it expects its opponent to make in response. Based on this search, and some clever heuristics that indicate what constitutes a good game position for the computer player, the best sequence of moves can be chosen. This searching technique relies on the fact that there are usually a relatively small number of moves that a player can make at

any one time in a game. However, the fact that the ancient Chinese game of Go-Moku has not, to date, been mastered by computer players (van der Werf et al. 2002) illustrates the restrictions of such techniques.

The common thread linking together the kinds of games that this article focuses on is that they all contain computer controlled virtual characters that possess a strong notion of agency. Efforts are often made to separate the many different kinds of modern video games that are the focus of this article into a small set of descriptive genres. Unfortunately, much like in music, film and literature, no categorization can hope to perfectly capture the nuances of all of the available titles. However, a brief mention of some of the more important game genres is worth while (a more detailed description of game genres, and artificial intelligence requirements of each is given in Laird and van Lent 2000).

The most popular game genre is without doubt the *action game* in which the player must defeat waves of demented foes, typically (for increasingly bizarre motivations) bent upon global destruction. Illustrative examples of the genre include *Half-Life 2* (www.half-life2.com) and the *Halo* series (www.halo3.com). A screenshot of the upcoming action game *Rogue Warrior* (www.bethsoft.com) is shown in Fig. 2.

Strategy *games* allow players to control large armies in battle with other people, or computer



Computer Graphics and Games, Agent-Based Modeling in, Fig. 2 A screenshot of the upcoming action game *Rogue Warrior* from Bethesda Softworks. (Image courtesy of Bethesda Softworks)



Computer Graphics and Games, Agent-Based Modeling in, Fig. 3 A screenshot from Bethesda Softwork's role playing game *The Elder Scrolls IV: Oblivion*. (Image courtesy of Bethesda Softworks)

controlled opponents. Players do not have direct control over their armies, but rather issue orders which are carried out by agent-based artificial soldiers. Well regarded examples of the genre include the *Age of Empires* (www.ageofempires.com) and *Command & Conquer* (www.commandandconquer.com) series.

Role playing games (such as the *Elder Scrolls* (www.elderscrolls.com) series) place game players in expansive virtual worlds across which they must embark on fantastical quests which typically involve a mixture of solving puzzles, fighting opponents and interacting with non-player characters in order to gain information. Figure 3 shows a screenshot of the aforementioned role-playing game *The Elder Scrolls IV: Oblivion*.

Almost every sport imaginable has at this stage been turned into a computer based *sports game*. The challenges in developing these games are creating computer controlled opponents and team mates that play the games at a level suitable to the human player. Interesting examples include *FIFA Soccer 08* (www.fifa08.ea.com) and *Forza Motorsport 2* (www.forzamotorsport.net).

Finally, many people expected that the rise of massively multi-player online games (MMOGs),

in which hundreds of human players can play together in an online world, would sound the death knell for the use of virtual non-player characters in games. Examples of MMOGs include *World of Warcraft* (www.worldofwarcraft.com) and *Battlefield 2142* (www.battlefield.ea.com). However, this has not turned out to be the case as there are still large numbers of single player games being produced and even MMOGs need computer controlled characters for roles that players do not wish to play.

Of course there are many games that simply do not fit into any of these categorizations, but that are still relevant for a discussion of the use of agent-based techniques – for example *The Sims* (www.thesims.ea.com) and the *Microsoft Flight Simulator* series (www.microsoft.com/games/flightsimulatorx). However the categorization still serves to introduce those unfamiliar with the subject to the kinds of games up for discussion.

Implementing Agent-Based Modelling Techniques in Games

One of the earliest examples of using agent-based modeling techniques in video games was its application to path planning. The ability of non-player

characters (NPCs) to manoeuvre around a game world is one of the most basic competencies required in games. While in very early games it was sufficient to have NPCs move along pre-scribed paths, this soon become unacceptable. Games programmers soon began to turn to AI techniques which might be applied to solve some of the problems that were arising. The A* path planning algorithm (Stout 1996) was the first example of such a technique to find wide-spread use in the games industry. Using the A* algorithm NPCs can be given the ability to find their own way around an environment. This was put to particularly fine effect early on in real-time strategy games where the units controlled by players are semi-autonomous and are given orders rather than directly controlled. In order to use the A* algorithm a game world must be divided into a series of cells each of which is given a rating in terms of the effort that must be expended to cross it. The A* algorithm then performs a search across these cells in order to find the shortest path that will take a game agent from a start position to a goal.


Since becoming widely understood amongst the game development community many interesting additions have been made to the basic A* algorithm. It was not long before three dimensional versions of the algorithm became commonly used (Smith 2002). The basic notion of storing the energy required to cross a cell within a game world has also been extended to augment cells with a wide range of other useful information (such as the level of danger in crossing a cell) that can be used in the search process (Reed and Geisler 2003).

The next advance in the kind of techniques being used to achieve agent-based modeling in games was the finite state machine (FSM) (Houlette and Fu 2003). An FSM is a simple system in which a finite number of *states* are connected in a directed graph by *transitions* between these states. When used for the control of NPCs, the nodes of an FSM indicate the possible actions within a game world that an agent can perform. Transitions indicate how changes in the state of the game world or the character's own attributes (such as health, tiredness etc) can move the agent from one state to another.

Figure 4 shows a sample FSM for the control of an NPC in a typical action game. In this example the behaviors of the character are determined by just four states – CHASE, ATTACK, FLEE and EXPLORE. Each of these states provides an action that the agent should take. For example, when in the EXPLORE state the character should wander randomly around the world, or while in the FLEE state the character should determine a direction to move in that will take it away from its current enemy and move in that direction. The links between the states show how the behaviors of the character should move between the various available states. So, for example, if while in the ATTACK state the agent's health measure becomes low, they will move to the FLEE state and run away from their enemy.

FSMs are widely used because they are so simple, well understood and extremely efficient both in terms of processing cycles required and memory usage. There have also been a number of highly successful augmentations to the basic state machine model to make them more effective, such as the introduction of layers of parallel state machines (Alexander 2003), the use of fuzzy logic in finite state machines (Dybsand 2001) and the implementation of cooperative group behaviors through state machines (Snaveley 2002).

The action game *Halo 2* is recognized as having a particularly good implementation of state



Computer Graphics and Games, Agent-Based Modeling in, Fig. 4 A simple finite state machine for a soldier NPC in an action game

machine based NPC control (Valdes 2004). At any time an agent could be in any one of the four states *Idle*, *Guard/Patrol*, *Attack/Defend*, and *Retreat*. Within each of these states a set of rules was used in order to select from a small set of appropriate actions for that state – for example a number of different ways to attack the player. The decisions made by NPCs were influenced by a number of character attributes including strength, speed and cowardliness. Transition between states was triggered by perceptions made by characters simulated senses of vision and hearing and by internal attributes such as health. The system implemented also allowed for group behaviors allowing NPCs to hold conversations and cooperate to drive vehicles.

However, FSMs are not without their drawbacks. When designing FSMs developers must envisage every possible situation that might confront an NPC over the course of a game. While this is quite possible for many games, if NPCs are required to move between many different situations this task can become overwhelming. Similarly, as more and more states are added to an FSM designing the links between these states can become a mammoth undertaking.

From IGDA (2003) the definition of rule based systems states that they are “... *comprised of a database of associated rules. Rules are conditional program statements with consequent actions that are performed if the specified conditions are satisfied*”. Rule based systems have been applied extensively to control NPCs in games (Christian 2002), in particular for the control of NPCs in role-playing games. NPCs behaviors are scripted using a set of rules which typically indicate how an NPC should respond to particular events within the game world. Borrowed from Woodcock (2000), the listing below shows a snippet of the rules used to control a warrior character in the RPG *Baldur's Gate* (www.bioware.com).

```
IF
// If my nearest enemy is not within 3
!Range(NearestEnemyOf(Myself),3)
// and is within 8
Range(NearestEnemyOf(Myself),8)
THEN
// 1/3 of the time
```

```
RESPONSE #40
// Equip my best melee weapon
EquipMostDamagingMelee()
// and attack my nearest enemy, checking every
60
// ticks to make sure he is still the nearest
AttackReevalutate(NearestEnemyOf
(Myself),60)
// 2/3 of the time
RESPONSE #60
// Equip a ranged weapon
EquipRanged()
// and attack my nearest enemy, checking every
30
// ticks to make sure he is still the nearest
AttackReevalutate(NearestEnemyOf (Myself),
30)
```

The implementation of an NPC using a rule-based system would consist of a large set of such rules, a small set of which would fire based on the conditions in the world at any given time. Rule based systems are favored by game developers as they are relatively simple to use and can be exhaustively tested. Rule based systems also have the advantage that rule sets can be written using simple proprietary scripting systems (Berger 2002), rather than full programming languages, making them easy to implement. Development companies have also gone so far as to make these scripting languages available to the general public, enabling them to author their own rule sets.

Rule based systems, however, are not without their drawbacks. Authoring extensive rule sets is not a trivial task, and they are *usually* restricted to simple situations. Also, rule based systems can be restrictive in that they don't allow sophisticated interplay between NPCs motivations, and require that rule set authors foresee every situation that the NPC might find itself in.

Some of the disadvantages of simple rule based systems can be alleviated by using more sophisticated inference engines. One example uses Dempster Schafer theory (Laramée 2002) which allows rules to be evaluated by combining multiple sources of (often incomplete) *evidence* to determine actions. This goes some way towards supporting the use of rule based systems in situations where complete knowledge is not available.

ALife techniques have also been applied extensively in the control of game NPCs, as much as a philosophy as any particular techniques. The outstanding example of this is *The Sims* (thesims.ea.com) a surprise hit of 2000 which has gone on to become the best selling PC game of all time. Created by games guru Will Wright the Sims puts the player in control of the lives of a virtual family in their virtual home. Inspired by aLife, the characters in the game have a set of motivations, such as hunger, fatigue and boredom and seek out items within the game world that can satisfy these desires. Virtual characters also develop sophisticated social relationships with each other based on common interest, attraction and the amount of time spent together. The original system in the Sims has gone on to be improved in the sequel *The Sims 2* and a series of expansion packs.

Some of the more interesting work in developing techniques for the control of game characters (particularly in action games) has been *focused* on developing interesting sensing and memory models for game characters. Players expect when playing action games that computer controlled opponents should suffer from the same problems that players do when perceiving the world. So, for example, computer controlled characters should not be able to see through walls or from one floor to the next. Similarly, though, players expect computer controlled characters to be capable of perceiving events that occur in a world and so NPCs should respond appropriately to sound events or on seeing the player.

One particularly fine example of a sensing model was in the game *Thief: The Dark Project* where players are required to sneak around an environment without alerting guards to their presence (Leonard 2003). The developers produced a relatively sophisticated sensing model that was used by non-player characters which modeled visual effects such as not being able to see the player if they were in shadows, and moving some way towards modeling acoustics so that non-player characters could respond reasonably to sound events.

2004s *Fable* (fable.lionhead.com) took the idea of adding memory to a game to new heights. In this adventure game the player took on the role of a

hero from boyhood to manhood. However, every action the player took had an impact on the way in which the *game* world's population would react to him or her as they would remember every action the next time they met the player. This notion of long-term consequences added an extra layer of believability to the game-playing experience.

Serious Games & Academia

It will probably have become apparent to most readers of the previous section that much of the work done in implementing agent-based techniques for the control of NPCs in commercial games is relatively simplistic when compared to the application of these techniques in other areas of more academic focus, such as robotics (Muller 1996). The reasons for this have been discussed already and briefly relate to the lack of available processing resources and the requirements of commercial quality control. However, a large amount of very interesting work is taking place in the application of agent-based technologies in academic research, and in particular the field of serious games. This section will begin by introducing the area of serious games and then go on to discuss interesting academic projects looking at agent-based technologies in games.

The term serious games (Michael and Chen 2005) refers to games designed to do more than just entertain. Rather, serious games, while having many features in common with conventional games, have ulterior motives such as teaching, training, and marketing. Although games have been used for ends apart from entertainment, in particular education, for a long time, the modern serious games movement is set apart from these by the level of sophistication of the games it creates. The current generation of serious games is comparable with main-stream games in terms of the quality of production and sophistication of their design. Serious games offer particularly interesting opportunities for the use of agent-based modeling techniques due to the facts that they often do not have to live up to the rigorous testing of commercial games, can have the requirement of specialized hardware rather than being restricted to commercial games hardware and often, by the nature of their application

domains, require more in-depth interactions between players and NPCs.

The modern serious games movement can be said to have begun with the release of *America's Army* (www.americasarmy.com) in 2002 (Nieborg 2004). Inspired by the realism of commercial games such as the *Rainbow 6* series (www.rainbow6.com), the United States military developed America's Army and released it free of charge in order to give potential recruits a flavor of army life. The game was hugely successful and is still being used today as both a recruitment tool and as an internal army training tool.

Spurred on by the success of *America's Army* the serious games movement began to grow, particularly within academia. A number of conferences sprung up and notably the Serious Games Summit became a part of the influential Game Developer's Conference (www.gdconf.com) in 2004.

Some other notable offerings in the serious games field include *Food Force* (www.foodforce.com) (DeMaria 2005), a game developed by the United Nations World Food Programme in order to promote awareness of the issues surrounding emergency food aid; *Hazmat Hotzone*

(Carless 2005), a game developed by the Entertainment Technology Centre at Carnegie Mellon University to train fire-fighters to deal with chemical and hazardous materials emergencies; *Yourself!Fitness* (www.yourselffitness.com) (Michael and Chen 2005) an interactive virtual personal trainer developed for modern games consoles; and *Serious Gordon* (www.seriousgames.ie) (Mac Namee et al. 2006) a game developed to aid in teaching food safety in kitchens. A screen shot of *Serious Gordon* is shown in Fig. 5.

Over the past decade, interest in academic research that is directly focused on artificial intelligence, and in particular agent-based modelling techniques and their application to games (as opposed to the general virtual character/computer graphics work discussed previously) has grown dramatically. One of the first major academic research projects into the area of Game-AI was led by John Laird at the University of Michigan, in the United States. The SOAR architecture was developed in the early nineteen eighties in an attempt to "develop and apply a unified theory of human and artificial intelligence" (Rosenbloom et al. 1993). SOAR is essentially a rule based



Computer Graphics and Games, Agent-Based Modeling in, Fig. 5 A screenshot of *Serious Gordon*, a serious game developed to aid in the teaching of food safety in kitchens

inference system which takes the current state of a problem and matches this to production rules which lead to actions.

After initial applications into the kind of simple puzzle worlds which characterized early AI research (Laird et al. 1984), the SOAR architecture was applied to the task of controlling computer generated forces (Jones et al. 1999). This work lead to an obvious transfer to the new research area of game-AI (Laird 2000).

Initially the work of Laird's group focused on applying the SOAR architecture to the task of controlling NPC opponents in the action game *Quake* (www.idsoftware.com) (Laird 2000). This proved quite successful leading to opponents which could successfully play against human players, and even begin to plan based on anticipation of what the player was about to do. More recently Laird's group have focused on the development of a game which requires more involved interactions between the player and the NPCs. Named *Haunt 2*, this game casts the player in the role of a ghost that must attempt to influence the actions of a group of computer controlled characters inhabiting the ghost's haunted house (Magerko

et al. 2004). The main issue that arises with the use the SOAR architecture is that it is enormously resource hungry, with the NPC controllers running on a separate machine to the actual game.

At Trinity College in Dublin in Ireland, the author of this article worked on an intelligent agent architecture, the Proactive Persistent Agent (PPA) architecture, for the control of background characters (or support characters) in character-centric games (games that focus on character interactions rather than action, e.g. role-playing games) (Mac Namee and Cunningham 2003; Mac Namee et al. 2003). The key contributions of this work were that it made possible the creation of NPCs that were capable of behaving believably in a wide range of situations and allowed for the creation of game environments which it appeared had an existence beyond their interactions with players. Agent behaviors in this work were based on models of personality, emotion, relationships to other characters and behavioral models that changed according to the current role of an agent. This system was used to develop a stand alone game and as part of a simulation of areas within Trinity College. A screenshot of this second application is shown in Fig. 6.



Computer Graphics and Games, Agent-Based Modeling in, Fig. 6 Screenshots of the PPA system simulating parts of a college

At Northwestern University in Chicago the Interactive Entertainment group has also applied approaches from more traditional research areas to the problems facing game-AI. Ian Horswill has led a team that are attempting to use architectures traditionally associated with robotics for the control of NPCs. In Horswill and Zubek (1999) consider how perfectly matched the behavior based architectures often used in robotics are with the requirements of NPC control architectures. The group have demonstrated some of their ideas in a test-bed environment built on top of the game Half-Life (Khoo and Zubek 2002). The group also looks at issues around character interaction (Zubek and Horswill 2005) and the many psychological issues associated with creating virtual characters asking how we can create virtual game agents that display all of the foibles that make us relate to characters in human stories (Horswill 2007).

Within the same research group a team led by Ken Forbus have extended research previously undertaken in conjunction with the military (Forbus et al. 1991) and applied it to the problem of terrain analysis in computer strategy games (Forbus et al. 2001). Their goal is to create strategic opponents which are capable of performing sophisticated reasoning about the terrain in a game world and using this knowledge to identify complex features such as ambush points. This kind of high level reasoning would allow AI opponents play a much more realistic game, and even surprise human players from time to time, something that is sorely missing from current strategy games.

As well as this work which has spring-boarded from existing applications, a number of projects began expressly to tackle problems in game-AI. Two which particularly stand out are the Excalibur Project, led by Alexander Nareyek (2001) and work by John Funge (1999). Both of these projects have attempted to applying sophisticated planning techniques to the control of game characters.


Nareyek uses constraint based planning to allow game agents reason about their world. By using techniques such as local search Nareyek has attempted to allow these sophisticated agents

perform resource intensive planning within the constraints of a typical computer game environment. Following on from this work, the term *anytime agent* was coined to describe the process by which agents actively refine original plans based on changing world conditions. In Nareyek (2007) describes the directions in which he intends to take this work in future.

Funge uses the situational calculus to allow agents reason about their world. Similarly to Nareyek he has addressed the problems of a dynamic, ever changing world, plan refining and incomplete information. Funge's work uses an extension to the situational calculus which allows the expression of uncertainty. Since completing this work Funge has gone on to be one of the founders of AiLive (www.aillive.net), a middleware company specializing in AI for games.

While the approaches of both of these projects have shown promise within the constrained environments to which they have been applied during research, (and work continues on them) it remains to be seen whether such techniques can be successfully applied to a commercial game environment and all of the resource constraints that such an environment entails.

One of the most interesting recent examples of agent-based work in the field of serious games is that undertaken by Barry Silverman and his group at the University of Pennsylvania in the United States (Silverman et al. 2006a, b). Silverman models the protagonists in military simulations for use in training programmes and has taken a very interesting approach in that his agent models are based on established cognitive science and behavioral science research. While Silverman admits that many of the models described in the cognitive science and behavioral science literature are not well quantified enough to be directly implemented, he has adapted a number of well respected models for his purposes. Silverman's work is an excellent example of the capabilities that can be explored in a serious games setting rather than a commercial game setting, and as such merits an in depth discussion. A high-level schematic diagram of Silverman's approach is shown in Fig. 7 and shows the agent architecture used by Silverman's system, PMFserv.



Computer Graphics and Games, Agent-Based Modeling in, Fig. 7 A schematic diagram of the main components of the PMFserv system. (With kind permission of Barry Silverman)

The first important component of the PMFserv system is the biology module which controls biological needs using a metaphor based on the flow of water through a system. Biological concepts such as hunger and fatigue are simulated using a series of reservoirs, tanks and valves which model the way in which resources are consumed by the system. This biological model is used in part to model stress which has an important impact on the way in which agents make decisions. To model the way in which agent performance changes under pressure Silverman uses *performance moderator functions* (PMFs). An example of one of the earliest PMFs used is the Yerkes–Dodson “inverted-u” curve (Yerkes and Dodson 1908) which illustrates that as mental arousal is increased performance initially improves, peaks and then trails off again. In PMFserv a range of PMFs are used to model the way in which behavior should change depending on stress levels and biological conditions.

The second important module of PMFserv attempts to model how personality, culture and emotion affect the behavior of an agent. In keeping with the rest of their system PMFserv uses models inspired by cognitive science to model

emotions. In this case the well known OCC model (Ortony et al. 1988), which has been used in agent-based applications before (Bates 1992b), is used. The OCC model provides for 11 pairs of opposite emotions such as pride and shame, and hope and fear. The emotional state of an agent with regard to past, current and future actions heavily influences the decisions that the agent makes.

The second portion of the Personality, Culture, Emotion module uses a value tree in order to capture the values of an agent. These values are divided into a Preference Tree which captures long term desired states for the world, a Standards Tree which relates to the actions that an agent believes it can or cannot follow in order to achieve these desired states and a Goal Tree which captures short term goals.

PMFserv also models the relationships between agents (Social Model, Relations, Trust in Fig. 7). The relationship of one agent to another is modeled in terms of three axes. The first is the degree to which the other agent is thought of as a human rather than an inanimate object – locals tend to view American soldiers as objects rather than people. The second axis is the cognitive



Computer Graphics and Games, Agent-Based Modeling in, Fig. 8 A screenshot of the PMFserv system being used to simulate the Black Hawk Down scenario. (With kind permission of Barry Silverman)

grouping (ally, foe etc) to which the other agent belongs and whether this is also a group to which the first agent has an affinity. Finally, the valence, or strength, of the relationship is stored. Relationships continually change based on actions that occur within the game world. Like the other modules of the system this model is also based on psychological research (Ortony et al. 1988).

The final important module of the PMFserv architecture is the Cognitive module which is used to decide on particular actions that agents will undertake. This module uses inputs from all of the other modules to make these decisions and so the behavior of PMFserv agents is driven by their stress levels, relationships to other agents and objects within the game world, personality, culture and emotions. The details of the PMFserv cognitive process are beyond the scope of this article, so it will suffice to say that action selection is based on a calculation of the utility of a particular action to an agent, with this calculation modified by the factors listed above.

The most highly developed example using the PMF-serv model is a simulation of the 1993 event

in Mogadishu, Somalia in which a United States military Black Hawk helicopter crashed, as made famous by the book and film “Black Hawk Down” (Bowden 2000). In this example, which was developed as a military training aid as part of a larger project looking at agent implementations within such systems (Toth et al. 2003; van Lent et al. 2004) the player took on the role of a US army ranger on a mission to secure the helicopter wreck in a modification (or “mod”) of the game *Unreal Tournament* (www.unreal.com). A screenshot of this simulation is shown in Fig. 8.

The PMFserv system was used to control the behaviors of characters within the game world such as Somali militia, and Somali civilians. These characters were imbued with physical attributes, a value system and relationships with other characters and objects within the game environment. The sophistication of PMFserv was apparent in many of the behaviors of the simulations NPCs. One particularly good example was the fact that Somali women would offer themselves as human shields for militia fighters. This behavior was never directly programmed into the agents

make-up, but rather emerged as a result of their values and assessment of their situation. PMFserv remains one of the most sophisticated current agent implementations and shows the possibilities when the shackles of commercial game constraints are thrown off.

Future Directions

There is no doubt that with the increase in the amount of work being focused on the use of agent-based modeling in computer graphics and games there will be major developments in the near future. This final section will attempt to predict what some of these might be.

The main development that might be expected in all of the areas that have been discussed in this article is an increase in the depth of simulation. The primary driver of this increase in depth will be the development of more sophisticated agent models which can be used to drive ever more sophisticated agent behavior. The PMFserv system described earlier is one example of the kinds of deeper systems that are currently being developed. In general computer graphics applications this will allow for the creation of more interesting simulations including previously prohibitive features such as automatic realistic facial expressions and other physical expressions of agents' internal states. This would be particularly use in CGI for movies in which, although agent based modeling techniques are commonly used for crowd scenes and background characters, main characters are still animated almost entirely by hand.

In the area of computer games it can be expected that many of the techniques being used in movie CGI will filter over to real-time game applications as the processing power of game hardware increases – this is a pattern that has been evident for the past number of years. In terms of depth that might be added to the control of game characters one feature that has mainly been conspicuous by its absence in modern games is genuine learning by game agents. 2000s *Black & White* and its sequel *Black & White 2* (www.lionhead.com) featured some learning by one of the game's main characters

that the player could teach in a reinforcement manner (Evans 2002). While this was particularly successful in the game, such techniques have not been more widely applied. One interesting academic project in this area is the NERO project (www.nerogame.org) which allows a player to train an evolving army of soldiers and have them battle the armies of other players (Stanley et al. 2006). It is expected that these kinds of capabilities will become more and more common in commercial games.

One new feature of the field of virtual character control in games is the emergence of specialized middleware. Middleware has had a massive impact in other areas of game development including character modeling (for example Maya available from www.autodesk.com) and physics modeling (for example Havok available from www.havok.com). AI focused middleware for games is now becoming more common with notable offerings including AI-Implant (www.ai-implant.com) and Kynogon (www.kynogon.com) which perform path finding and state machine based control of characters. It is expected that more sophisticated techniques will over time find their way into such software.

To conclude the great hope for the future is that more and more sophisticated agent-based modeling techniques from other application areas and other branches of AI will find their way into the control of virtual characters.

Bibliography

Primary Literature

- Adamson A (Director) (2005) The chronicles of Narnia: the Lion, the Witch and the Wardrobe. Motion Picture. http://adisney.go.com/disneypictures/narnia/lb_main.html
- Aitken M, Butler G, Lemmon D, Saindon E, Peters D, Williams G (2004) The Lord of the Rings: the visual effects that brought middle earth to the screen. In: International conference on computer graphics and interactive techniques (SIGGRAPH), course notes
- Alexander T (2003) Parallel-state machines for believable characters. In: Massively multiplayer game development. Charles River Media
- Allers R, Minkoff R (Directors) (1994) The Lion King. Motion picture. <http://disney.go.com/disneyvideos/animatedfilms/lionking/>

- Aylett R, Luck M (2000) Applying artificial intelligence to virtual reality: intelligent virtual environments. *Appl Artif Intell* 14(1):3–32
- Badler N, Bindiganavale R, Bourne J, Allbeck J, Shi J, Palmer M (1999) Real Time virtual humans. In: Proceedings of the international conference on digital media futures
- Bates J (1992a) The nature of characters in interactive worlds and the Oz project. Technical report CMU-CS-92–200. School of Computer Science, Carnegie Mellon University
- Bates J (1992b) Virtual reality, art, and entertainment. *Presence J Teleoper Virtual Environ* 1(1):133–138
- Berger L (2002) Scripting: overview and code generation. In: Rabin S (ed) *AI game programming wisdom*. Charles River Media, Hingham
- Bird B, Pinkava J (Directors) (2007) *Ratatouille*. Motion picture. <http://disney.go.com/disneyvideos/animatedfilms/ratatouille/>
- Blumberg B (1996) Old tricks, new dogs: ethology and interactive creatures. Ph.D. Thesis, Media Lab, Massachusetts Institute of Technology
- Bowden M (2000) *Black Hawk Down*. Corgi Adult
- Burke R, Isla D, Downie M, Ivanov Y, Blumberg B (2002) Creature smarts: the art and architecture of a virtual brain. In: Proceedings of game-on 2002: the 3rd international conference on intelligent games and simulation, pp 89–93
- Burton T (Director) (1992) *Batman returns*. Motion picture. <http://www.warnervideo.com/batmanmoviesondvd/>
- Carless S (2005) Postcard from SGS 2005: Hazmat: hot-zone – first-person first responder gaming. Retrieved Oct 2007, from Gamasutra: www.gamasutra.com/features/20051102/carless_01b.shtml
- Christian M (2002) A simple inference engine for a rule based architecture. In: Rabin S (ed) *AI game programming wisdom*. Charles River Media, Hingham
- Darnell E, Johnson T (Directors) (1998) *Antz*. Motion picture. <http://www.dreamworksanimation.com/>
- DeMaria R (2005) Postcard from the serious games summit: how the United Nations fights hunger with food force. Retrieved Oct 2007, from Gamasutra: www.gamasutra.com/features/20051104/demaria_01.shtml
- Dybsand E (2001) A generic fuzzy state machine in C++. In: Rabin S (ed) *Game programming gems 2*. Charles River Media, Hingham
- Evans R (2002) Varieties of learning. In: Rabin S (ed) *AI game programming wisdom*. Charles River Media, Hingham
- Faloutsos P, van de Panne M, Terzopoulos D (2001) The virtual stuntman: dynamic characters with a repertoire of autonomous motor skills. *Comput Graph* 25(6):933–953
- Farenc N, Musse S, Schweiss E, Kallmann M, Aune O, Boulic R et al (2000) A paradigm for controlling virtual humans in urban environment simulations. *Appl Artif Intell J Special Issue Intell Virtual Environ* 14(1):69–91
- Feng-Hsiung H (2002) Behind deep blue: building the computer that defeated the world chess champion. Princeton University Press, Princeton
- Forbus K, Nielsen P, Faltings B (1991) Qualitative spatial reasoning: the CLOCK project. *Artif Intell* 51:1–3
- Forbus K, Mahoney J, Dill K (2001) How qualitative spatial reasoning can improve strategy game AIs. In: Proceedings of the AAAI spring symposium on AI and interactive entertainment
- Funge J (1999) *AI for games and animation: a cognitive modeling approach*. A.K. Peters, Natick
- Hayes-Roth B, Doyle P (1998) Animate characters. *Auton Agents Multi-Agent Syst* 1(2):195–230
- Horswill I (2007) Psychopathology, narrative, and cognitive architecture (or: why NPCs should be just as screwed up as we are). In: Proceedings of AAAI fall symposium on intelligent narrative technologies
- Horswill I, Zubek R (1999) Robot architectures for believable game agents. In: Proceedings of the 1999 AAAI spring symposium on artificial intelligence and computer games
- Houlette R, Fu D (2003) The ultimate guide to FSMs in games. In: Rabin S (ed) *AI game programming wisdom 2*. Charles River Media, Hingham
- IGDA (2003) Working group on rule-based systems report. International Games Development Association
- Isbister K, Doyle P (2002) Design and evaluation of embodied conversational agents: a proposed taxonomy. In: Proceedings of the AA-MAS02 workshop on embodied conversational agents: lets specify and compare them! Bologna
- Jackson P (Director) (2001) *The lord of the rings: the fellowship of the ring*. Motion picture. <http://www.lordoftherings.net/>
- Jackson P (Director) (2002) *The lord of the rings: the two towers*. Motion picture. <http://www.lordoftherings.net/>
- Jackson P (Director) (2003) *The lord of the rings: the return of the king*. Motion picture. <http://www.lordoftherings.net/>
- Johnston O, Thomas F (1995) *The illusion of life: Disney animation*. Disney Editions, New York
- Jones R, Laird J, Neilsen P, Coulter K, Kenny P, Koss F (1999) Automated intelligent pilots for combat flight simulation. *AI Mag* 20(1):27–42
- Khoo A, Zubek R (2002) Applying inexpensive AI techniques to computer games. *IEE Intell Syst Spec Issue Interact Entertain* 17(4):48–53
- Koeppel D (2002) Massive attack. <http://www.popsci.com/popsci/science/d726359b9fa84010vgncvm100004eeCBCdr.crd.html>. Accessed Oct 2007
- Laird J (2000) An exploration into computer games and computer generated forces. In: The 8th conference on computer generated forces and behavior representation
- Laird J, van Lent M (2000) Human-level AI's killer application: interactive computer games. In: Proceedings of the 17th national conference on artificial intelligence
- Laird J, Rosenbloom P, Newell A (1984) Towards Chunking as a general learning mechanism. In: The 1984 national conference on artificial intelligence (AAAI), pp 188–192
- Laramée F (2002) A rule based architecture using Dempster-Schafer theory. In: Rabin S (ed) *AI game programming wisdom*. Charles River Media, Hingham

- Lasseeter J, Stanton A (Directors) (1998) *A Bug's life*; Motion picture. <http://www.pixar.com/featurefilms/abl/>
- Leonard T (2003) Building an AI sensory system: examining the design of thief: the dark project. In: Proceedings of the 2003 game developers' conference, San Jose
- Loyall B (1997) *Believable agents: building interactive personalities*. Ph.D. thesis, Carnegie Melon University
- Määta A (2002) Realistic level design for Max Payne. In: Proceedings of the 2002 game developer's conference, GDC 2002
- Mac Namee B, Cunningham P (2003) Creating socially interactive non player characters: the μ -SIC system. *Int J Intell Games Simul* 2(1)
- Mac Namee B, Dobbyn S, Cunningham P, O'Sullivan C (2003) Simulating virtual humans across diverse situations. In: Proceedings of intelligent virtual agents'03, pp 159–163
- Mac Namee B, Rooney P, Lindstrom P, Ritchie A, Boylan F, Burke G (2006) Serious gordon: using serious games to teach food safety in the kitchen. In: The 9th international conference on computer games: AI, animation, mobile, educational & serious games CGAMES06, Dublin
- Magerko B, Laird JE, Assanie M, Kerfoot A, Stokes D (2004) AI characters and directors for interactive computer games. In: The 2004 innovative applications of artificial intelligence conference. AAAI Press, San Jose
- Michael D, Chen S (2005) *Serious games: games that educate, train, and inform*. Course Technology PTR
- Muller J (1996) *The design of intelligent agents: a layered approach*. Springer, Berlin
- Musse RS, Thalmann D (2001) A behavioral model for real time simulation of virtual human crowds. *IEEE Trans Vis Comput Graph* 7(2):152–164
- Nareyek A (2001) *Constraint based agents*. Springer, Berlin
- Nareyek A (2007) Game AI is dead. Long live game AI! *IEEE Intell Syst* 22(1):9–11
- Nieborg D (2004) America's army: more than a game. Bridging the gap; Transforming knowledge into action through gaming and simulation. In: Proceedings of the 35th conference of the international simulation and gaming association (ISAGA), Munich
- Ortony A, Clore GL, Collins A (1988) *The cognitive structure of emotions*. Cambridge University Press, Cambridge
- Perlin K, Goldberg A (1996) Improv: a system for scripting interactive actors in virtual worlds. In: Proceedings of the ACM computer graphics annual conference, pp 205–216
- Proyas A (Director) (2004) *I, Robot*. Motion picture. <http://www.irobotmovie.com>
- Rao AS, Georgeff MP (1991) Modeling rational agents within a BDI-architecture. In: Proceedings of knowledge representation and reasoning (KR&R-91). Morgan Kaufmann, pp 473–484
- Reed C, Geisler B (2003) Jumping, climbing, and tactical reasoning: how to get more out of a navigation system. In: Rabin S (ed) *AI game programming wisdom* 2. Charles River Media, Hingham
- Reynolds C (1987) Flocks, herds and schools: a distributed behavioral model. *Comput Graph* 21(4):25–34
- Rodriguez R (Director) (1996) *From Dusk 'Till Dawn*. Motion picture
- Rosenbloom P, Laird J, Newell A (1993) *The SOAR papers: readings on integrated intelligence*. MIT Press
- Sánchez-Crespo D (2006) *GDC: physical gameplay in Half-Life 2*. Retrieved Oct 2007, from gamasutra.com: http://www.gamasutra.com/features/20060329/sanchez_01.shtml
- Shao W, Terzopoulos D (2005) Autonomous pedestrians. In: Proceedings of SIGGRAPH/EG symposium on computer animation, SCA'05, pp 19–28
- Silverman BG, Bharathy G, O'Brien K, Cornwell J (2006a) Human behavior models for agents in simulators and games: part II: Gamebot engineering with PMFserv. *Presence Teleoper Virtual Worlds* 15(2):163–185
- Silverman BG, Johns M, Cornwell J, O'Brien K (2006b) Human behavior models for agents in simulators and games: part I: enabling science with PMFserv. *Presence Teleoper Virtual Environ* 15(2):139–162
- Smith P (2002) Polygon soup for the programmer's soul: 3D path finding. In: Proceedings of the game developer's conference 2002, GDC2002
- Snively P (2002) Agent cooperation in FSMs for baseball. In: Rabin S (ed) *AI game programming wisdom*. Charles River Media, Hingham
- Stanley KO, Bryant BD, Karpov I, Miikkulainen R (2006) RealTime evolution of neural networks in the NERO video game. In: Proceedings of the twenty-first national conference on artificial intelligence, AAAI-2006. AAAI Press, pp 1671–1674
- Stout B (1996) Smart moves: intelligent path-finding. *Game Dev Mag* Oct
- Takahashi TS (1992) Behavior simulation by network model. *Mem Kougakuin Univ* 73, pp 213–220
- Terzopoulos D, Tu X, Grzeszczuk R (1994) Artificial fishes with autonomous locomotion, perception, behavior and learning, in a physical world. In: Proceedings of the artificial life IV workshop. MIT Press
- Thalmann MN, Thalmann D (1994) *Artificial life and virtual reality*. Wiley, Chichester
- Thompson C (2007) *Halo 3: how Microsoft labs invented a new science of play*. Retrieved Oct 2007, from wired.com: http://www.wired.com/gaming/virtualworlds/magazine/15-09/ff_halo
- Toth J, Graham N, van Lent M (2003) Leveraging gaming in DOD modelling and simulation: integrating performance and behavior moderator functions into a general cognitive architecture of playing and non-playing characters. In: Twelfth conference on behavior representation in modeling and simulation (BRIMS, formerly CGF), Scotsdale
- Valdes R (2004) *In the mind of the enemy: the artificial intelligence of Halo 2*. Retrieved Oct 2007, from How-StuffWorks.com: <http://entertainment.howstuffworks.com/halo2-ai.htm>

- van der Werf E, Uiterwijk J, van den Herik J (2002) Programming a computer to play and solve Ponnuki-go. In: Proceedings of game-on 2002: the 3rd international conference on intelligent games and simulation, pp 173–177
- van Lent M, McAlinden R, Brobst P (2004) Enhancing the behavioral fidelity of synthetic entities with human behavior models. In: Thirteenth conference on behavior representation in modeling and simulation (BRIMS)
- Woodcock S (2000) AI roundtable moderator's report. In: Proceedings of the game developer's conference 2000 (GDC2000)
- Wooldridge M, Jennings N (1995) Intelligent agents: theory and practice. *Know Eng Rev* 10(2):115–152
- Yerkes RW, Dodson JD (1908) The relation of strength of stimulus to rapidity of habit formation. *J Comp Neurol Psychol* 18:459–482
- Zubek R, Horswill I (2005) Hierarchical parallel Markov models of interaction. In: Proceedings of the artificial intelligence and interactive digital entertainment conference, AIIDE 2005

Books and Reviews

- DeLoura M (ed) (2000) *Game programming gems*. Charles River Media, Hingham
- DeLoura M (ed) (2001) *Game programming gems 2*. Charles River Media, Hingham
- Dickheiser M (ed) (2006) *Game programming gems 6*. Charles River Media, Hingham
- Kirmse A (ed) (2004) *Game programming gems 4*. Charles River Media, Hingham
- Pallister K (ed) (2005) *Game programming gems 5*. Charles River Media, Hingham
- Rabin S (ed) (2002) *Game AI wisdom*. Charles River Media, Boston
- Rabin S (ed) (2003) *Game AI wisdom 2*. Charles River Media, Boston
- Rabin S (ed) (2006) *Game AI wisdom 3*. Charles River Media, Boston
- Russell S, Norvig P (2002) *Artificial intelligence: a modern approach*. Prentice Hall
- Treglia D (ed) (2002) *Game programming gems 3*. Charles River Media, Hingham

Agent-Based Modeling, Large-Scale Simulations

Hazel R. Parry

Central Science Laboratory, York, UK

Article Outline

Glossary

Definition of the Subject

Introduction

Large Scale Agent Based Models: Guidelines for
Development

Parallel Computing

Example

Future Directions

Bibliography

Glossary

Agent A popular definition of an agent, particularly in AI research, is that of Wooldridge (Wooldridge 1999), pp. 29: “an agent is a computer system that is situated in some environment, and that is capable of autonomous action in this environment in order to meet its design objectives”. In particular, it is the autonomy, flexibility, inter-agent communication, reactivity and proactive-ness of the agents that distinguishes the paradigm and gives power to agent-based models and multi-agent simulation (Heppenstall 2004; Jennings 2000). Multi-agent systems (MAS) comprise of numerous agents, which are given rules by which they act and interact with one another to achieve a set of goals.

Block Mapping A method of partitioning an array of elements between nodes of a distributed system, where the array elements are partitioned as evenly as possible into blocks of consecutive elements and assigned to processors. The size of the blocks approximates to the number of array elements divided by the number of processors.

Complexity Complexity and complex systems pertain to ideas of randomness and irregularity in a system, where individual-scale interactions may result in either very complex or surprisingly simple patterns of behavior at the larger scale. Complex agent-based systems are therefore usually made up of agents interacting in a non-linear fashion. The agents are capable of generating emergent behavioral patterns, of deciding between rules and of relying upon data across a variety of scales. The concept allows for studies of interaction between hierarchical levels rather than fixed levels of analysis.

Cyclic Mapping A method of partitioning an array of elements between nodes of a distributed system, where the array elements are partitioned by cycling through each node and assigning individual elements of the array to each node in turn.

Grid Computer ‘Grids’ are comprised of a large number of disparate computers (often desktop PCs) that are treated as a virtual cluster when linked to one another via a distributed communication infrastructure (such as the internet or an intranet). Grids facilitate sharing of computing, application, data and storage resources. Grid computing crosses geographic and institutional boundaries, lacks central control, and is dynamic as nodes are added or removed in an uncoordinated manner. BOINC computing is a form of distributed computing is where idle time on CPUs may be used to process information (<http://boinc.berkeley.edu/>).

Ising-type Model Ising-type models have been primarily used in the physical sciences. They simulate behavior in which individual elements (e.g., atoms, animals, social behavior, etc.) modify their behavior so as to conform to the behavior of other individuals in their vicinity. Conway’s Game of Life is a Ising-type model, where cells are in one of two states: dead or alive. In biology, the technique is used to model neural networks and flocking birds, for example.

Message Passing (MP) Message passing (MP) is the principle way by which parallel clusters of machines are programmed. It is a widely-used, powerful and general method of enabling distribution and creating efficient programs (Pacheco 1997). Key advantages of using MP architectures are an ability to scale to many processors, flexibility, ‘future-proofing’ of programs and portability (Openshaw and Turton 2000).

Message Passing Interface (MPI) A computing standard that is used for programming parallel systems. It is implemented as a library of code that may be used to enable message passing in a parallel computing system. Such libraries have largely been developed in C and Fortran, but are also used with other languages such as Java (MPIJava <http://www.hpjava.org>). It enables developers of parallel software to write parallel programs that are both portable and efficient.

Multiple Instruction Multiple Data (MIMD)

Parallelization where different algorithms are applied to different data items on different processors.

Parallel Computer Architecture A parallel computer architecture consists of a number of identical units that contain CPUs (Central Processing Units) which function as ordinary serial computers. These units, called nodes, are connected to one another (Fig. 1). They may transfer information and data between one another (e.g. via MPI) and simultaneously perform calculations on different data.

Single Instruction Multiple Data (SIMD)

SIMD techniques exploit data level parallelism: when a large mass of data of a uniform

type needs the same instruction performed on it. An example is a vector or array processor. An application that may take advantage of SIMD is one where the same value is being added (or subtracted) to a large number of data points.


Vector Computer/Vector Processor Vector computers contain a CPU designed to run mathematical operations on multiple data elements simultaneously (rather than sequentially). This form of processing is essentially a SIMD approach. The Cray Y-MP and the Convex C3880 are two examples of vector processors used for supercomputing in the 1980s and 1990s. Today, most recent commodity CPU designs include some vector processing instructions.

Definition of the Subject

‘Large scale’ simulations in the context of agent-based modelling are not only simulations that are large in terms of the size of the simulation (number of agents simulated), but they are also complex. Complexity is inherent in agent-based models, as they are usually composed of dynamic, heterogeneous, interacting agents. Large scale agent-based models have also been referred to as ‘Massively Multi-agent Systems (MMAS)’ (Ishida et al. 2005). MMAS is defined as “‘beyond resource limitation’: the number of agents exceeds local computer resources, or the situations are too complex to design/program given human cognitive resource limits’ (Ishida et al. 2005), Preface. Therefore, for agent-based modelling ‘large scale’ is not simply a size problem, it is

Agent-Based Modeling, Large-Scale Simulations,

Fig. 1 A network with interconnected separate memory and processors. (After Pacheco 1997, pp. 19)



also a problem of managing complexity to ensure scalability of the agent model.

Multi-agent simulation models increase in scale as the modeller requires many agents to investigate whole system behavior, or the modeller wishes to fully examine the response of a single agent in a realistic context. Two key problems may be introduced by increasing the scale of a multi-agent system: (1) Computational resources limit the simulation time and/or data storage capacity and (2) Agent model analysis may become more difficult. Difficulty in analyzing the model may be due to the model system having a large number of complex components or to memory for model output storage being restricted by computer resources.

Introduction

Many systems that are now simulated using agent-based models are systems where agent numbers are large and potentially complex. These large scale models are constructed under a number of diverse scientific disciplines, with differing histories of agent simulation and methodologies emerging with which to deal with large scale simulations: for example molecular physics, social science (e.g. crowd simulation, city growth), telecommunications, ecology and military research. The primary

methodology to emerge is parallel computing, where an agent model is distributed across a number of CPUs to increase the memory and processing power available to the simulation. However, there are a range of potential methods with which to simulate large numbers of agents. Some suggestions are listed in Table 1.

The simplest solution to enable larger scale agent simulation is usually to improve the computer hardware that is used and run the model on a server or invest in a more powerful PC. However, this option may be too costly or may not provide enough scaling so other options may then be considered. Another simple solutions may be to reduce the number of agents or to revert to a simpler modelling approach such as a population model, but both of these solutions would significantly alter the model and the philosophy behind the model, probably not addressing the research question that the model was initially constructed for. There are a number of unique advantages and insights that may be gained from an agent-based approach. Agent simulations are often constructed to enable the analysis of emergent spatial patterns and individual life histories. The realism of an agent-based approach may be lost by using a simpler modelling technique.

The structure of agent simulations (often with asynchronous updating and heterogeneous data

Agent-Based Modeling, Large-Scale Simulations, Table 1 Potential solutions to implement when faced with a large number of agents to model

Solution	Pro	Con
Reduce the number of agents in order for model to run	No reprogramming of model	Unrealistic population. Alters model behavior
Revert to a population based modelling approach	Could potentially handle any number of individuals	Lose insights from agent approach. Unsuitable for research questions. Construction of entirely new model (non-agent-based)
Invest in an extremely powerful computer	No reprogramming of model	High cost
Run the model on a vector computer	Potentially more efficient as more calculations may be performed in a given time	This approach only works more efficiently with SIMD, probably unsuitable for agent-based models
Super-individuals (Scheffer et al. 1995)	Relatively simple solution, keeping model formulation the same	Reprogramming of model. Inappropriate in a spatial context (Parry 2006; Parry and Evans In press)
Invest in a powerful computer network and reprogram the model in parallel	Makes available high levels of memory and processing power	High cost Advanced computing skills required for restructuring of model

types) would mean that running a simulation on a vector computer may make little difference to the simulation performance. This is because an agent model typically has few elements that could take advantage of SIMD: rarely the same value will be added (or subtracted) to a large number of data points. Vector processors are less successful when a program does not have a regular structure, and they don't scale to arbitrarily large problems (the upper limit on the speed of a vector program will be some multiple of the speed of the CPU (Pacheco 1997)).

Another relatively simple option is to implement an aggregation of the individual agents into 'super-agents', such as the 'super-individual' approach of Scheffer et al. (1995). The basic concept of this approach is shown in Fig. 2. These 'super-agents' are formed from individual agents that share the same characteristics, such as age and sex. However, it may not be possible to group agents in a simulation in this way and, importantly, this method has been proven ineffective in a spatial context (Parry 2006; Parry and Evans *In press*).


The most challenging solution, to reprogram the model in parallel, is a popular solution due to the shortcomings of the other approaches outlined above. A parallel solution may also have some monetary cost and require advanced computing skills to implement, but it can potentially greatly increase the scale of the agent simulation. Extremely large scale object-oriented simulations that simulate

individual particles on massively parallel computer systems have been successfully developed in the physical sciences of fluid dynamics, meteorology and materials science. In the early 1990s, work in the field of molecular-dynamics (MD) simulations proved parallel platforms to be highly successful in enabling large-scale MD simulation of up to 131 million particles (Lomdahl et al. 1993). Today the same code has been tested and used to simulate up to 320 billion atoms on the BlueGene/L architecture containing 131,072 IBM PowerPC440 processors (Kadau et al. 2006). These simulations include calculations based upon the short-range interaction between the individual atoms, thus in some ways approximate to agent simulations although in other ways molecular-dynamic simulations lack the complexity of most agent-based models (see Table 2).

There are significant decisions to be made when considering the application of a computing solution such as parallel programming to solve the problem of large numbers of agents. In addition to the issue of reprogramming the model to run on a parallel computer architecture, it is also necessary to consider the additional complexity of agents (as opposed to atoms), so that programming models and tools facilitate the deployment, management and control of agents in the distributed simulation (Gasser et al. 2005). For example, distributed execution resources and timelines must be managed, full encapsulation of agents must be enforced, and tight control over message-based

Agent-Based Modeling, Large-Scale Simulations,

Fig. 2 'Super-agents':
Grouping of individuals
into single objects that
represent the collective



Agent-Based Modeling, Large-Scale Simulations, Table 2 Key elements of a ‘bottom-up’ simulation that may affect the way in which it may scale. Agent

simulations tend to be complex (to the right of the table), though may have some elements that are less complex, such as local or fixed interactions

Element	Least complex	Most complex
Spatial structure	Aspatial or Lattice of cells (1d, 2d or 3d +)	Continuous space
Internal state	Simple representation (boolean true or false)	Complex representation (many states from an enumerable set) or fuzzy variable values
Agent heterogeneity	No	Yes
Interactions	Local and fixed (within a neighborhood)	Multiple different ranges and stochastic
Synchrony of model updates	Synchronous update	Not synchronous: asynchrony due to state-transition rules

multi-agent interactions is necessary (Gasser et al. 2005). Agent models can vary in complexity, but most tend to be complex especially in the key model elements of spatial structure and agent heterogeneity. Table 2 gives an indication of the relative complexity of model elements found in models that focus on individual interactions (which encompasses both multi-agent models and less complex, ‘Ising’-type models).

The following sections detail guidelines for the development of a large scale agent-based model, highlighting in particular the challenges faced in writing large scale, high performance Agent based Modelling (ABM) simulations and giving a suggested development protocol. Following this, an example is given of the parallelization of a simple agent-based model, showing some of the advantages but also the pitfalls of this most popular solution. Key challenges, including difficulties that may arise in the analysis of agent-based models at a large scale, are highlighted. Alternative solutions are then discussed and some conclusions are drawn on the way in which large scale agent-based simulation may develop in coming years.

Large Scale Agent Based Models: Guidelines for Development

Key Considerations

There is no such thing as a standard agent-based model, or even a coherent methodology for agent simulation development (although recent

literature in a number of fields sets out some design protocols, e.g. Gilbert (2007) and Grimm et al. (2006)). Thus, there can be no standard method to develop a large scale agent-based model. However, there are certain things to consider when planning to scale up a model. Some key questions to ask about the model are as follows:

1. What program design do you already have and what is the limitation of this design?
 - (a) What is it the memory footprint for any existing implementation?
 - (b) What are your current run times?
2. What are your scaling requirements?
 - (a) How much do you need to scale now?
 - (b) How far do you need to scale eventually?
 - (c) How soon do you need to do it?
3. How simple is your model and how is it structured?
4. What are your agent complexities?
5. What are your output requirements?

The first question is to identify the limitations in the program design that you are using and to focus on the primary ‘bottlenecks’ in the model. These limitations will either be due to memory or speed (or perhaps both). Therefore it will be necessary to identify the memory footprint for your existing model, and analyze run times, identifying where the most time is taken or memory used by the simulation. It is primarily processor power that controls the speed of the simulation. Runtime will

also increase massively once Random Access Memory (RAM) is used up, as most operating systems will resort to virtual memory (i. e. hard drive space), thus a slower mechanism with mechanical parts rather than solid-state technology engages. At this stage, it may be such that simple adjustments to the code may improve the scalability of the model. However, if the code is efficient, other solutions will then need to be sought.

The second question is how much scaling is actually necessary for the model. It may be such that a simple or interim solution (e.g. upgrading computer hardware) may be acceptable whilst only moderate scaling is required, but longer term requirements should also be considered – a hardware upgrade may be a quick fix but if the model may eventually be used for much larger simulations it is necessary to plan for the largest scaling that will potentially be required.

The third question, relating to model simplicity and structure, is key to deciding a methodology that can be used to scale a model up. A number of factors will affect whether a model will be easy to distribute in parallel, for example. These include whether the model iterates at each time step or is event driven, whether it is aspatial or spatial and the level/type of agent interaction (both with one another and with the environment). More detail on the implications of these factors is given in section “[Parallel Computing](#)”.

Agent complexity, in addition to model structure, may limit the options available for scaling up a model. For example, a possible scaling solution may be to group individual agents together as ‘super-individuals’ (Scheffer et al. 1995). However, if agents are too complex it may not be possible to determine a simple grouping system (such as by age), as agent behavior may be influenced heavily by numerous other state variables.

Output requirements are also important to consider. These may already be limiting the model, in terms of memory for data storage. Even if they are not currently limiting the model in this way, once the model is scaled up output data storage needs may be an issue, for example, if the histories of individual agents need to be stored. In addition, the way that output data is handled by the model may be altered if the model structure is altered

(e.g. if agents are grouped together output will be at an aggregate level). Thus, an important consideration is to ensure that output data is comparable to the original model and that it is feasible to output once the model structure is altered.

A Protocol

In relation to the key considerations highlighted above, a simple protocol for developing a large scale agent-based simulation can be defined as follows:

1. Optimize existing code.
2. Clearly identify scaling requirements (both for now and in the future).
3. Consider simple solutions first (e.g. a hardware upgrade).
4. Consider more challenging solutions.
5. Evaluate the suitability of the chosen scaling solution on a simplified version of the model before implementing on the full model.

The main scaling solution to implement (e.g. from Table 1) is defined by the requirements of the model. Implementation of more challenging solutions should be done in stages, where perhaps a simplified version of the model is implemented on a larger scale. Agent simulation development should originate with a local, flexible ‘prototype’ and then as the model development progresses and stabilizes larger scale implementations can be experimented with Gasser et al. (2005). This is necessary for a parallel implementation of a model, for example, as a simplified model enables and assessment of whether it is likely to provide the desired improvements in model efficiency. This is particularly the case for improvements in model speed, as this depends on improved processing performance that is not easily calculated in advance.

Parallel Computing

Increasing the capacity of an individual computer in terms of memory and processing power has limited ability to perform large scale agent simulations, particularly due to the time the machine

would take to run the model using a single processor. However, by using multiple processors and a mix of distributed and shared memory working simultaneously, the scale of the problem for each individual computer is much reduced. Subsequently, simulations can run in a fraction of the time that would be taken to perform the same complex, memory intensive, operations. This is the essence of parallel computing. ‘Parallel computing’ encompasses a wide range of computer architectures, from a HPC (High performance computing) Linux box, to dedicated multi-processor/multi-core systems (such as a Beowulf cluster), super clusters, local computer clusters or Grids and public computing facilities (e.g. Grid computers, such as the White Rose Grid, UK <http://www.wrgrid.org.uk/>). The common factor is that these systems consist of a number of interconnected ‘nodes’ (processing units), that may perform simultaneous calculations on different data. These calculations may be the same or different, depending whether a ‘Single Instruction Multiple Data’ (SIMD) or ‘Multiple Instruction Multiple data’ (MIMD) approach is implemented.

In terms of MAS, parallel computing has been used to develop large scale agent simulations in a number of disciplines. These range from ecology, e.g. Abbott et al. (1997), Immanuel et al. (2005), Lorek and Sonnenschein (1995), and Wang et al. (2004, 2005b, 2006a, 2006b) and biology, e.g. Castiglione et al. (1997) and Da-Jun et al. (2004) to social science, e.g. Takeuchi (2005) and computer science, e.g. Popov et al. (2003), including artificial intelligence and robotics, e.g. Bokma et al. (1994) and Bouzid et al. (2001).

Several key challenges arise when implementing an agent model in parallel, which may affect the increase in performance achieved. These include load balancing between nodes, synchronizing events to ensure causality, monitoring of the distributed simulation state, managing communication between nodes and dynamic resource allocation (Timm and Pawlaszczyk 2005). Good load balancing and inter-node communication with event synchronisation are central to the development of an efficient parallel simulation, and are further discussed below.

Load Balancing

In order to ensure the most efficient use of memory and processing resources in a parallel computing system the data load must be balanced between processors and the work load equally distributed. If this is not the case then one computer may be idle as others are working, resulting in time delays and inefficient use of the system’s capacity. There are a number of ways in which data can be ‘mapped’ to different nodes and the most appropriate depends on the model structure. Further details and examples are given in Pacheco (1997), including ‘block mapping’ and ‘cyclic mapping’. An example of ‘block mapping’ load balancing is given below, in section “[Example](#)”.

In many simulations the computational demands on the nodes may alter over time, as the intensity of the agents’ or environment’s processing requirements varies on each node over time. In this case, dynamic load balancing techniques can be adopted to further improve the parallel model performance. For example, Jang (2006) and Jang and Agha (2005), use a form of dynamic load balancing with object migration they term “Adaptive Actor Architecture”. Each agent platform monitors the workload of its computer node and the communication patterns of agents executing on it in order to redistribute agents according to their communication localities as agent platforms become overloaded. However, this approach does introduce additional processing overheads, so is only worth implementing for large scale agent simulations where some agents communicate with one another more intensely than other agents (communication locality is important) or communication patterns are continuously changing so static agent allocation is not efficient.

Communication Between Nodes

It is important to minimize inter-node communication when constructing a parallel agent simulation, as this may slow the simulation down significantly if the programmer is not careful (Takahashi and Mizuta 2006; Takeuchi 2005). The structure of the model itself largely determines the way in which data should be split and information transferred between nodes to

maximize efficiency. Agent simulations generally by definition act spatially within an environment. Thus, an important first consideration is whether to split the agents or the environment between nodes. The decision as to whether to split the agents between processors or elements of the environment such as grid cells largely depends upon the complexity of the environment, the mobility of the agents, and the number of interactions between the agents. If the environment is relatively simple (thus information on the whole environment may be stored on all nodes), it is probably most efficient to distribute the agents. This is particularly the case if the agents are highly mobile, as a key problem when dividing the environment between processors is the transfer of agents or information between processors. However, if there are complex, spatially defined interactions between agents, splitting agents between nodes may be problematic, as agents may be interacting with other agents that are spatially local in the context of the whole simulation but are residing on different processors. Therefore conversely, if the agents are not very mobile but have complex, local interactions and/or the agents reside in a complex environment, it is probably best to split the environment between nodes (Logan and Theodoropolous 2001). Further efficiency may be achieved by clustering the agents which communicate heavily with each other (Takahashi and Mizuta 2006).

In models where there is high mobility and high interaction it is often possible, especially for ecological models, to find a statistical commonality that can be used as a replacement for more detailed interaction. For example, as will be shown in our example, if the number of local agent interactions is the only important aspect of the interactions, a density map of the agents, transferred to a central node, aggregated and redistributed, might allow agents to be divided between nodes without the issue of having to do detailed inter-agent communication between nodes a large number of times.

The way in which the simulation iterates may influence the approach taken when parallelizing the model. The model may update synchronously at a given time step or asynchronously (usually

because the system is event-driven). In addition, agents may update asynchronously but the nodes may be synchronized at each time step or key model stage. Asynchronous updating may be a problem if there is communication between nodes, as some nodes may have to wait for others to finish processes before communication takes place and further processing is possible, resulting in blocking (see below). Communication between nodes then becomes highly complex (Wang et al. 2005a). It is important that messages communicating between agents are received in the correct order, however a common problem in distributed simulations is ensuring that this is so as other factors, such as latency in message transmission across the network, may affect communication (Wang et al. 2005a).

A number of time management mechanisms exist that may be implemented to manage message passing in order to ensure effective node to node communication, e.g. Fujimoto (1998).

Blocking and Deadlocking

Deadlock occurs when two or more processes are waiting for communication from one of the other processes. When programming a parallel simulation it is important to avoid deadlock to ensure the simulation completes. The simplest example is when two processors are programmed to receive from the other processor before that processor has sent. This may be simply resolved by changing the order that tasks are executed, or to use ‘non-blocking’ message passing. Where blocking is used, processing on nodes waits until a message is transmitted. However, when ‘non-blocking’ is used, processing continues even if the message hasn’t been transmitted yet. The use of a non-blocking MPI may reduce computing times, and work can be performed while communication is in progress.

Example

To demonstrate some of the benefits and pitfalls of parallel programming for a large scale agent-based model, a simple example is given here. This summarizes a simplified agent-based model

of aphid population dynamics in agricultural landscapes of the UK, which was parallelized to cope with millions of agents, as described in detail in Parry (2006), Parry and Evans (In press) and Parry et al. (2006).

A key problem with the original, non-parallel, aphid simulation was that it was hindered by memory requirements, which were far larger than could be accommodated at any individual processing element. This is a common computing problem (Chalmers and Tidmus 1996). The data storage required for each aphid object in a landscape scale simulation quickly exceeded the storage capacity of a PC with up to 2097 MB of RAM. The combined or ‘virtual shared’ memory of several computers was used to cope with the amount of data needed, using a Single Instruction Multiple-Data approach (SIMD).

Message-passing techniques were used to transfer information between processors, to distribute the agents in the simulation across a Beowulf cluster (a 30-node distributed memory parallel computer). A Message-passing Interface (MPI) for Java was used, MPIJava (<http://www.hpjava.org>). ‘MPIJava wraps around the open-source’ open-source native MPI ‘LAM’ (<http://www.lam-mpi.org/>). Further details on the methods used to incorporate the MPI into the model are given in Parry (2006) and Parry et al. (2006).

Effective parallelization minimizes the passing of information between nodes, as it is processor intensive. In the example model, only the environment object and information on the number of agents to create on each node are passed from a single control node to each of the other nodes in the cluster, and only density information is returned to the control node for redistribution and display. The control node manages the progress of the model, acts as a central communication point for the model and handles any code that may not be distributed to all nodes (such as libraries from an agent toolkit or a GUI). Structuring a model without a control node is possible, or the control node may also be used to process data, depending on the requirements of the simulation. Transfer of density values, rather than agents, significantly reduced the computational

overheads for message passing between the nodes. The model was simple enough that specific interagent communication between nodes was not necessary.

Even distribution of data between nodes was achieved by splitting immigrant agents evenly across the system, with each node containing information on the environment and local densities passed from the control node. The number of immigrants to be added to each node was calculated by a form of ‘block mapping’, pp. 35 in Pacheco (1997), which partitioned the number of immigrants into blocks which were then assigned to each node. So, if there were three nodes ($n \text{ D } 3$) and thirteen immigrants ($i \text{ D } 13$), the immigrants mapped to each node would be as follows:

$$i_0, i_1, i_2, i_3 \rightarrow n_1$$

$$i_4, i_5, i_6, i_7 \rightarrow n_2$$

$$i_8, i_9, i_{10}, i_{11}, i_{12} \rightarrow n_3.$$

As thirteen does not divide evenly by three, the thirteenth agent is added to the final node.

Benefits


Simulation runtime and memory availability was greatly improved by implementing the simple aphid model in parallel across a large number of nodes. The greatest improvement in simulation runtime and memory availability was seen when the simulation was run across the maximum number of nodes (25) (Figs. 3 and 4). The largest improvement in speed given by the parallel model in comparison to the non-parallel model is when more than 500,000 agents are run across 25 nodes, although the parallel model is slower by comparison for lower numbers. This means that additional processing power is required in the parallel simulation compared to the original model, such that only when very large numbers of agents are run does it become more efficient.

Pitfalls

Although there are clear benefits of distributing the example simulation across a large number of nodes, the results highlight that the parallel


Agent-Based Modeling, Large-Scale Simulations,

Fig. 3 Plot of the mean maximum memory used (per node) against number of agents for the model: comparison between simulations using 2, 5 and 25 nodes and the non-parallel model (single processor)



Agent-Based Modeling, Large-Scale Simulations,

Fig. 4 Plot of the percentage speed up (per node) from the non-parallel model against number of agents for the model: comparison between simulations using 2, 5 and 25 nodes and the non-parallel model (single processor)



approach is not always more efficient than the original single processor implementation of a model. In the example, the two node simulation used more memory on the worker node than the non-parallel model when the simulation had 100,000 agents or above. This is due to additional memory requirements introduced by message passing and extra calculations required in the parallel implementation (which are less significant when more nodes are used as these requirements remain relatively constant).

The results also highlight that adding more processors does not necessarily increase the model speed. The example model shows that for simulations run on two nodes (one control node, one worker node) the simulation takes longer to run in parallel compared to the non-parallel model. Message passing time delay and the modified structure of the code are responsible. The

greatest improvement in speed is when more than 500,000 agents are run across 25 nodes, however when lower numbers of nodes are used the relationship between the number of nodes and speed is complex: for 100,000 agents five nodes are faster than the non-parallel model, but for 500,000 the non-parallel model is faster. Overall, these results suggest that when memory is sufficient on a single processor, it is unlikely to ever be efficient to parallelize the code, as when the number of individuals was low the parallel simulation took longer and was less efficient than the non-parallel model run on a single node.

This demonstrates that in order to effectively parallelize an agent model, the balance between the advantage of increasing the memory availability and the cost of communication between nodes must be assessed. By following an iterative development process as suggested in section

“A Protocol”, the threshold below which parallelization is not efficient and whether this option is suitable for the model should become apparent in the early stages of model development. Here, the simplified study confirmed the value of further development to build a parallel version of the full model.

Future Directions

Complexity and Model Analysis

In addition to the processing and data handling issues faced by large scale agent-based simulations, as agent simulations increase in scale and complexity, model analysis may become more difficult and the system may become intractable. There is no clearly defined way of dealing with increased difficulties of model analysis introduced by the greater complexity of large scale agent-based models, in the same way that there is no clear way to deal with the complexity inherent in most agent-based systems. However, some guidelines to model analysis have recently been developed for agent simulations, e.g. Grimm et al. (2006), and some suggestions are put forward on ways in which agent simulation complexity may be described, for example by the use of a model ‘overview’, ‘design concepts’ and ‘details’ (ODD) protocol for agent model description. In particular, the protocol requires that design concepts are linked to general concepts identified in the field of Complex Adaptive Systems (Grimm and Railsback 2005), including emergence and stochasticity.

To address issues of complexity in a large scale agent simulation, one possibility is to include an additional layer of structuring, ‘organizational design’, into the system. This is where the peer group of an agent, its roles and responsibilities are assigned in the model and made explicit, pp. 121 in Horling and Lesser (2005). In general agent simulations already have an implicit organization structure; Horling and Lesser (2005) argue that explicit organizational design highlights hidden inefficiencies in the model and allows the model to take full advantages of the resources available.

Grid Computing

For large scale agent-based models, many researchers have reached the limitations of ordinary PCs. However, there are ‘critical tensions’ (Gasser et al. 2005) between agent simulations built on ordinary PCs and heterogeneous, distributed parallel programming approaches. The architecture of a distributed system is very different to that of an ordinary PC, thus to transfer a simulation to a computer cluster additional system properties must be taken into account, including management of the distribution of the simulation and concurrency (Gasser et al. 2005). This is particularly apparent when parallelization is attempted on heterogeneous, non-dedicated systems such as a public Grid. The Grid may offer a number of advantages for large scale agent-based simulation, such as collaboration between modellers, access to resources and geographically distributed datasets (Zhang et al. 2005). However, in such systems, issues of infrastructure reliability, functional completeness and the state of documentation for some kinds of environments exist (Gasser et al. 2005). In order to use such a system the ‘fundamental’ issue of partial failures must be addressed (e.g. with a dynamic agent replication strategy (Guessoum et al. 2005)).

Dissemination of Techniques

For parallel computing as a solution to large scale agent-based simulation, there is an interesting and useful future challenge to develop user friendly, high performance, versatile hardware architectures and software systems. Many developers of agent simulations are not computer scientists by training, and still rely upon numerous agent toolkits for simulation development (e.g. Swarm and Repast). Automatic distribution of agents to whatever resources are available would be a great tool for many agent software developers. Therefore, perhaps the greatest challenge would be to develop a system that would allow for parallelisation to be performed in an agent simulation automatically, where agents may be written in a high-level language and could be automatically partitioned to nodes in a network. One example of an attempt to achieve this is Graphcode (<http://parallel.hpc.unsw.edu.au/rks/graphcode/>). Based

upon MPI, it maps agent-based models onto parallel computers, where agents are written based upon their graph topography to minimize communication overhead. Another example is HLA_GRID_Repast (Zhang et al. 2005), ‘a system for executing large scale distributed simulations of agent-based systems over the Grid’, for users of the popular Repast agent toolkit. HLA_GRID_Repast is a middleware layer which enables the execution of a federation of multiple interacting instances of Repast models across a grid network with a High Level Architecture (HLA). This is a ‘centralized coordination approach’ to distributing an agent simulation across a network (Timm and Pawlaszczyk 2005). Examples of algorithms designed to enable dynamic distribution of agent simulations are given in Scheutz and Schermerhorn (Scheutz and Schermerhorn 2006).

Although parallel computing is often the most effective way of handling large scale agent-based simulations, there are still some significant obstacles to the use of parallel computing for MAS. As shown with the simple example given here, this may not always be the most effective solution depending upon the increase in scale needed and the model complexity. Other possible methods were suggested in section “[Introduction](#)”, but these may also be unsuitable. Another option could be to deconstruct the model and simplify only certain elements of the model using either parallel computing or one of the other solutions suggested in section “[Introduction](#)”. Such a ‘hybrid’ approach is demonstrated by Zhang and Lui (2005), who combine equation-based approaches and multi-agent simulation with a Cellular Automata to simulate the complex interactions in the process of human immune response to HIV. The result is a model where equations are used to represent within-site processes of HIV infection, and agent-based simulation is used to represent the diffusion of the virus between sites.

It is therefore important to consider primarily the various ways in which the model may be altered, hybridized or simplified yet still address the core research questions, before investing money in hardware or investing time in the development of complex computational solutions.

Making the transition from a serial application to a parallel version is a process that requires a fair degree of formalism and program restructuring, so is not to be entered into lightly without exploring the other options and the needs of the simulation first.

Overall, it is clear that disparate work is being done in a number of disciplines to facilitate large scale agent-based simulation, and knowledge is developing rapidly. Some of this work is innovative and highly advanced, yet inaccessible to researchers in other disciplines who may unaware of key developments outside of their field. This chapter synthesizes and evaluates large scale agent simulation to date, providing a reference for a wide range of agent simulation developers.

Acknowledgments Many thanks to Andrew Evans (Multi-Agent Systems and Simulation Research Group, University of Leeds, UK) and Phil Northing (Central Science Laboratory, UK) for their advice on this chapter.

Bibliography

Primary Literature

- Abbott CA, Berry MW, Comiskey EJ, Gross LJ, Luh H-K (1997) Parallel individual-based modeling of everglades deer ecology. *IEEE Comput Sci Eng* 4(4):60–78
- Anderson J (2000) A generic distributed simulation system for intelligent agent design and evaluation. In: *Proceedings of the AI, simulation & planning in high autonomy systems*, Arizona
- Bokma A, Slade A, Kerridge S, Johnson K (1994) Engineering large-scale agent-based systems with consensus. *Robot Comput-Integr Manuf* 11(2):81–91
- Bouazid M, Chevrier V, Vialle S, Charpillet F (2001) Parallel simulation of a stochastic agent/environment interaction model. *Integr Comput-Aided Eng* 8(3):189–203
- Castiglione F, Bernaschi M, Succi S (1997) Simulating the immune response on a distributed parallel computer. *Int J Mod Phys C* 8(3):527–545
- Chalmers A, Tidmus J (1996) *Practical parallel processing: an introduction to problem solving in parallel*. International Thomson Computer Press, London
- Da-Jun T, Tang F, Lee TA, Sarda D, Krishnan A, Goryachev A (2004) Parallel computing platform for the agent-based modeling of multicellular biological systems. In: *Parallel and distributed computing: applications and technologies*. Lecture notes in computer science, vol 3320, pp 5–8
- Fujimoto RM (1998) Time management in the high level architecture. *Simulation* 71:388–400

- Gasser L, Kakugawa K (2002) MACE3J: fast flexible distributed simulation of large, large-grain multi-agent systems. In: *Proceedings of AAMAS*
- Gasser L, Kakugawa K, Chee B, Esteve M (2005) Smooth scaling ahead: progressive MAS simulation from single PCs to Grids. In: Davidsson P, Logan B, Takadama K (eds) *Multi-agent and multi-agent-based simulation. Joint Workshop MABS 2004*, New York, 19 July 2004. Springer, Berlin
- Gilbert N (2007) *Agent-based models. Quantitative applications in the social sciences*. SAGE, London
- Grimm V, Railsback SF (2005) *Individual-based modeling and ecology. Princeton series in theoretical and computational biology*. Princeton University Press, Princeton, 480 pp
- Grimm V, Berger U, Bastiansen F, Eliassen S, Ginot V, Giske J, Goss-Custard J, Grand T, Heinz S, Huse G, Huth A, Jepsen JU, Jorgensen C, Mooij WM, Muller B, Pe'er G, Piou C, Rails-back SF, Robbins AM, Robbins MM, Rossmannith E, Ruger N, Strand E, Souissi S, Stillman RA, Vabo R, Visser U, DeAngelis DL (2006) A standard protocol for describing individual-based and agent-based models. *Ecol Model* 198(1–2):115–126
- Guessoum Z, Briot J-P, Faci N (2005) Towards fault-tolerant massively multiagent system. In: Ishida T, Gasser L, Nakashima H (eds) *Massively multi-agent systems I: first international workshop MMAS 2004*, Kyoto, Dec 2004. Springer, Berlin
- Heppenstall AJ (2004) *Application of hybrid intelligent agents to modelling a dynamic, locally interacting retail market*. Ph D thesis, University of Leeds
- Horling B, Lesser V (2005) Quantitative organizational models for large-scale agent systems. In: Ishida T, Gasser L, Nakashima H (eds) *Massively multi-agent systems I: first international workshop MMAS 2004*, Kyoto, Japan, Dec 2004. Springer, Berlin
- Immanuel A, Berry MW, Gross LJ, Palmer M, Wang D (2005) A parallel implementation of ALFISH: simulating hydrological compartmentalization effects on fish dynamics in the Florida Everglades. *Simul Model Pract Theory* 13:55–76
- Ishida T, Gasser L, Nakashima H (eds) (2005) *Massively multi-agent systems I. First international workshop, MMAS 2004*, Kyoto. Springer, Berlin
- Jang MW (2006) Agent framework services to reduce agent communication overhead in large-scale agent-based simulations. *Simul Model Pract Theory* 14(6):679–694
- Jang MW, Agha G (2005) Adaptive agent allocation for massively multi-agent applications. In: Ishida T, Gasser L, Nakashima H (eds) *Massively multi-agent systems I: first international workshop MMAS 2004*, Kyoto, Dec 2004. Springer, Berlin
- Jennings NR (2000) On agent-based software engineering. *Artif Intell* 117:277–296
- Kadau K, Germann TC, Lomdahl PS (2006) Molecular dynamics comes of age: 320 billion atom simulation on BlueGene/L. *Int J Mod Phys C* 17(12):1755
- Lees M, Logan B, Theodoropoulos G (2002) Simulating agent-based systems with HLA: the case of SIM_AGENT. In: *Proceedings of the 2002 European simulation interoperability workshop*, pp 285–293
- Lees M, Logan B, Oguara T, Theodoropoulos G (2003) Simulating agent-based systems with HLA: the case of SIM_AGENT – Part II. In: *Proceedings of the 2003 European simulation interoperability workshop*
- Logan B, Theodoropolous G (2001) The distributed simulation of multi-agent systems. *Proc IEEE* 89(2):174–185
- Lomdahl PS, Beazley DM, Tamayo P, Gronbechjensen N (1993) Multimillion particle molecular-dynamics on the CM-5. *Int J Mod Phys C Phys Comput* 4(6):1075–1084
- Lorek H, Sonnenschein M (1995) Using parallel computers to simulate individual-oriented models in ecology: a case study. In: *Proceedings: ESM'95 European simulation multicongress*, Prague, June 1995
- Luke S, Cioffi-Revilla C, Panait L, Sullivan K (2004) MASON: a new multi-agent simulation toolkit. In: *Proceedings of the 2004 SwarmFest workshop*
- Openshaw S, Turton I (2000) *High performance computing and the art of parallel programming: an introduction for geographers, social scientists, engineers*. Routledge, London
- Pacheco PS (1997) *Parallel programming with MPI*. Morgan Kaufman Publishers, San Francisco
- Parry HR (2006) *Effects of land management upon species population dynamics: a spatially explicit, individual-based model*. Ph D thesis, University of Leeds
- Parry HR, Evans AJ (In press) A comparative analysis of parallel processing and super-individual methods for improving the computational performance of a large individual-based model. *Ecol Model*
- Parry HR, Evans AJ, Heppenstall AJ (2006) Millions of agents: parallel simulations with the repast agent-based toolkit. In: Trapp R (ed) *Cybernetics and systems 2006*, *Proceedings of the 18th European meeting on cybernetics and systems research*
- Popov K, Vlassov V, Rafea M, Holmgren F, Brand P, Haridi S (2003) Parallel agent-based simulation on a cluster of workstations. In: *EURO-PAR 2003 parallel processing*, vol 2790, pp 470–480
- Scheffer M, Baveco JM, DeAngelis DL, Rose KA, van Nes EH (1995) Super-individuals: a simple solution for modelling large populations on an individual basis. *Ecol Model* 80:161–170
- Scheutz M, Schermerhorn P (2006) Adaptive algorithms for the dynamic distribution and parallel execution of agent-based models. *J Parallel Distrib Comput* 66(8):1037–1051
- Takahashi T, Mizuta H (2006) Efficient agent-based simulation framework for multi-node supercomputers. In: Perrone LF, Wieland FP, Liu J, Lawson BG, Nicol DM, Fujimoto RM (eds) *Proceedings of the 2006 winter simulation conference*
- Takeuchi I (2005) A massively multi-agent simulation system for disaster mitigation. In: Ishida T, Gasser L, Nakashima H (eds) *Massively multi-agent systems I:*

- first international workshop MMAS 2004, Kyoto, Dec 2004. Springer, Berlin
- Timm IJ, Pawlaszczyk D (2005) Large scale multiagent simulation on the grid. In: Veit D, Schnizler B, Eymann T (eds) Proceedings of the workshop on agent-based grid economics (AGE 2005) at the IEEE international symposium on cluster computing and the grid (CCGRID). Cardiff University, Cardiff
- Wang D, Gross L, Carr E, Berry M (2004) Design and implementation of a parallel fish model for South Florida. In: Proceedings of the 37th annual Hawaii international conference on system sciences (HICSS'04)
- Wang F, Turner SJ, Wang L (2005a) Agent communication in distributed simulations. In: Davidsson P, Logan B, Takadama K (eds) Multi-agent and multi-agent-based simulation. Joint workshop MABS 2004, New York, 19 July 2004. Springer, Berlin
- Wang D, Carr E, Gross LJ, Berry MW (2005b) Toward ecosystem modeling on computing grids. *Comput Sci Eng* 7:44–52
- Wang D, Berry MW, Carr EA, Gross LJ (2006a) A parallel fish landscape model for ecosystem modeling. *Simulation* 82(7):451–465
- Wang D, Berry MW, Gross LJ (2006b) On parallelization of a spatially-explicit structured ecological model for integrated ecosystem simulation. *Int J High Perform Comput Appl* 20(4):571–581
- Wooldridge M (1999) Intelligent agents. In: Weiss G (ed) Multiagent systems: a modern approach to distributed artificial intelligence. MIT Press, Cambridge, pp 27–78
- Zhang S, Lui J (2005) A massively multi-agent system for discovering HIV-immune interaction dynamics. In: Ishida T, Gasser L, Nakashima H (eds) Massively multi-agent systems I: first international workshop MMAS 2004, Kyoto, Dec 2004. Springer, Berlin
- Zhang Y, Theodoropoulos G, Minson R, Turner SJ, Cai W, Xie Y, Logan B (2005) Grid-aware Large Scale Distributed Simulation of Agent-based Systems. In: 2004 European simulation interoperability workshop (EuroSIW 2005), 05E-SIW-047, Toulouse
- distributed network of agents over an MPI-based computer cluster.
- Graphcode system. <http://parallel.hpc.unsw.edu.au/rks/graphcode/>.
- MACE3J: <http://www.isrl.uiuc.edu/amag/mace/an> experimental agent platform supporting deployment of agent simulations across a variety of system architectures (Gasser and Kakugawa 2002; Gasser et al. (2005).
- MASON. <http://cs.gmu.edu/~eclab/projects/mason/>. MASON was 'not intended to include parallelization of a single simulation across multiple networked processors' (Luke et al. 2004). However, it does provide two kinds of simple parallelization:
1. Any given step in the simulation can be broken into parallel sub-steps each performed simultaneously.
 2. A simulation step can run asynchronously in the background independent of the simulation.
- Repast. <http://repast.sourceforge.net> and HLA_GRID_Repast (Zhang et al. 2005). The Repast toolkit has in-built capabilities for performing batch simulation runs.
- <http://www.cs.bham.ac.uk/research/projects/poplog/packages/simagent.html>. Two developments support distributed versions of SimAgent:
1. The use of HLA to distribute SimAgent (Lees et al. 2002, 2003)
 2. The SWAGES package: <http://www.nd.edu/~airolab/software/index.html>. This allows SimAgent to be distributed over different computers and interfaced with other packages.
- Message Passing Interfaces (MPI):
Background and tutorials. <http://www-unix.mcs.anl.gov/mpi/>
- MPICH2. <http://www-unix.mcs.anl.gov/mpi/mpich/>
- MPI forum. <http://www.mpi-forum.org/>
- MPIJava. <http://www.hpjava.org>
- OpenMP. <http://www.openmp.org>
- OpenMPI. <http://www.open-mpi.org/>
- Parallel computing and distributed agent simulation websites:
- Further references and websites <http://www.cs.rit.edu/~ncs/parallel.html>.
- Introduction to parallel programming. http://www.mhpcc.edu/training/workshop/parallel_intro/MAIN.html
- <http://www.agents.cs.nott.ac.uk/research/simulation/simulators/> (implementations of HLA_GRID_Repast and distributed SimAgent).
- Globus Grid computing resources. <http://www.globus.org/>.
- Beowulf computer clusters. <http://www.beowulf.org/>

Books and Reviews

- Agent libraries and toolkits with distributed or parallel features:
- Distributed GenSim: Supports distributed parallel execution (Anderson 2000).
- Ecolab: <http://ecolab.sourceforge.net/>. EcoLab models may also use the Graphcode library to implement a

Index

A

- Abreu-Rubinstein approach, 643
- Absorbing state, 785
- Abstract game, 9, 11, 624–626
 - network formation
 - stable sets in, 11
- Achievable mate, 371
- Acidosis, 857–858
- Acid-resistant phenotype, 859
- Action type, 185
- Active walker models, 669
- Adaptation, 725
- Adaptive behavior
 - behavior and cognition, 749
 - complex, 749
 - vs. concurrent developmental processes, 763
 - coupled dynamical processes, 753–755
 - definition, 747
 - developmental robotics methods, 762–763
 - as emergent dynamical properties, 749–753
 - evolutionary robotics methods, 761–762
 - incremental nature, developmental process, 763
 - methods, 760–761
 - multilevel and multi-scale organization, 755–759
 - social nature, developmental processes, 763
 - top-down effect, 759–760
- Additively separable preferences, 371
- Additivity, 428
- Adverse selection (hidden information), 287
- Agent (software agent), 667, 705, 778, 819, 865, 913
 - definition, 661
- Agent architecture, 667
- Agent-based computational economics (ACE), 662
 - agents, 829
 - Arthur, Holland, Lebaron, Palmer, and Tayler (AHLPT) stock market model, 828–829
 - autocorrelation of returns, 842
 - definition, 825–826
 - dividend changes, 829
 - efficient market believers (EMB), 835–836
 - excess volatility, 842
 - full spectrum, EMB investors, 844–845
 - heavy trading volume, 841–842
 - Kim and Markowitz portfolio insurers model, 827–828
 - LLS model, 830–835
 - Lux and Marchesi model, 829–830
 - Lux model, 829–830
 - market clearance, 836–837
 - no excess volatility, 838–841
 - number of rules, 829
 - return autocorrelation, 845
 - Schelling's segregation model, 826–827
 - two-EMB-population model, 843
- Agent-based modeling (ABM), 661, 663, 705, 725, 737, 865
 - agent behavior specification, 670–672
 - agent interaction, 674
 - ALife, 662
 - assumption, 869
 - autonomy, 676
 - CA-like models, 662
 - in computer generated imagery (CGI), 895–897
 - in computer graphics technology, 892–895
 - definition, 661, 667–668, 705, 727
 - direct interaction models, 675
 - discrete-event setup, 662
 - in games, 897–899
 - genotype/phenotype distinction, 731
 - idea of, 661
 - indirect interaction models, 676–677
 - large-scale, 663
 - methodology in science, 665
 - morphogenesis, 731–732
 - ontological approach, 676
 - ontology, 661
 - platforms for, 676–679
 - requirements, 868
 - and simulation, 663
- Agent-based simulation, 683, 786, 825
 - agent behavior specification, 670–672
 - agent interaction, 674
 - autonomy, 676
 - definition, 667–668, 683
 - direct interaction models, 675
 - as finite dynamical systems, 692–693
 - indirect interaction models, 676–677
 - ontological approach, 676
 - platforms for, 676–679

- Agent communication language (ACL)-based models, 674
- Aggregation rules, 511, 516, 518
- Allowable set of partners, 371
- All-pay auction, 340–341
- All-pay first-price auction, 335, 336, 340–341
- Anaerobic metabolism, 858
- Annotations, 865
- Antagonist games/zero-sum games, 306
- Ant colony optimization (ACO), 725, 737, 791, 801–802
- Approachability
 - Blackwell, 167
 - vs. excludability, 166
 - necessary and sufficient conditions, 165–166
 - no-regret strategies, 167–169
 - weak, 167
- Approximate cores, 476–477
- Arrow's desirable properties, 320–321
- Arrow's impossibility theorem, 511–514, 517, 529, 534–535
- Arthur, Holland, Lebaron, Palmer, and Tayler (AHLPT) stock market model, 828–829
- Artificial chemistry, 725
- Artificial intelligence (AI) techniques, 663, 868
- Artificial life (ALife), 662, 665, 725
 - in agent-based modeling, 738–743
 - algorithms, 734–735, 737
 - cellular automata, 729
 - and computing, 734
 - definition, 727
 - disciplines, 728
 - essence of, 728–730
 - game of life, 730–731
 - interaction, 738
 - population, 738
 - self-replication, 729
 - self-reproduction and replacement, 738
 - software programs, 728
 - sustainment and renewal, 738
 - Von Neumann's self-replication problem, 729
- Artificial neural network (ANN), 725, 736
- Artificial worlds, 769, 778–779
- Aspects, 865
- Assignment games, 17–20
- Asymmetric information, 287
- Asymptotic negligibility, 463
- Asynchronous swarms
 - definition, 807
 - local synchrony and self-synchronization, 809
 - modeling asynchrony, 807–809
 - natural, 809
 - realization, 810–811
 - types of, 807
- Attraction, singleton basins, 632–634
- Attributes, 865
- Auctions, 335
 - collusive bidding, 346–347
 - definition, 335
 - empirics of, 343–344
 - first-price sealed-bid, 338
 - outcomes, 341, 343
 - pre-auction knockout, 347
 - second-price, 337–338
 - single-item auction rules, 335–336
 - winner's curse, 344–346
- Aumann and Maschler model, 160
 - cav μ theorem, 162
 - concavification, 162
 - martingale, 161
 - recursive formula, 164
 - splitting, 161
- Aumann model, 122
- Aumann-Shapley rule, 451
- Autocatalytic set, 725, 733
- Autocorrelation of returns, 842
- Automata, 639
- Autonomous, 725
- Autonomy, 667
- Avida, 725
- Axiomatics, 56
 - characterizations, 54
- B**
- Babbling equilibrium, 251
- Bala-Goyal enforceability, 38
- Bala-Goyal rules, 611
- Bayesian equilibrium, 119, 132
 - and correlated equilibrium, 134–136
- Bayesian game, 107, 119, 131–132
 - correlated equilibrium in, 114
 - definition, 119
- Bayesian learning, 485
- Bayesian mechanism design, 327
- Bayesian-Nash equilibrium, 317, 335, 337–338
- Bayesian P¹-demand games, 449
- Behavioral type, 251
- Behavior strategy, 139, 251, 485
- Belief learning, 485, 487
- Belief subspaces, 125–129
- Bertrand duopoly, 33–34
- Bidding ring, 346
- Bidding strategy, 335, 337, 338, 345
- Bimatrix game, 171
- Biologically inspired computational algorithm, 725
- Blackboard-based architectures, 677
- Blackwell approachability, 167, 178
- Blackwell property, 165
- Blocking, 920
- Block mapping, 913
- Boltzmann equation, 775
- Borel's Poker model, 221–223
- Boundedness of marginal contributions (BMC), 478
- Bounded rationality, 825
- Brauer algebra, 718
- Brownian motion, 181
- Budget balance
 - condition, 424
 - cost sharing, 459
 - mechanism, 328

- Business languages, 875–876
- Butterfly effect, 769
- Bytecode, 865
- C**
- C#, 865, 887
- C++, 865, 887
- Cadherins, 851, 853
- Cake-cutting problem, 500
 - cut-and-choose, 501
 - dividing indivisible items, 503–506
 - picking sequences, 507
 - Selfridge-Conway procedure, 502
 - sequential algorithm (SA), 506
 - singles-doubles procedure (SD), 506
 - undercut procedure, 507
- Cancer
 - cell phenotypes, 852
 - progression, 852
- Capital accumulation, 230–231
- Cardinal utility, 438
- Cavu theorem, 162–164
- Cellular automata (CA), 662, 688, 725, 773, 778, 781, 791, 819
 - approaches, 669
 - cell-cell and cell-ECM adhesion, 856
 - classes, 731
 - definition, 851
 - discrete nature, 854
 - emergence of tumor invasion, 858–860
 - fiber tracts on glioma invasion, 857
 - game theory, 860
 - glioma cell motion, 857
 - heterogeneous environments on tumor cell dispersal, 857
 - invasive tumor morphology, 855
 - lattice-gas cellular automata, 855
 - metabolism and acidosis, 857–858
 - model, 854
 - notion of, 854
 - probabilistic, 784
 - tumor cell migration and extracellular matrix, 856–857
 - tumor invasion, 855–856
 - 4-tuple L,S,N,F definition, 854
- Cellular-computing architecture, 791
- Cellular-computing methods (CCm), 799
- Cellular Potts model, 856
- Chaos, 769
- Chaos theorems, 511, 517
- Characteristic/coalitional function, 49
- Characteristic function form game, 9
- Characterization, 67, 511
- Cheap-talk game, 251, 257–261
- Chemotaxis, 851, 855
- C-ImmSim, 687
- Class, 865
- Climate change, 61
- Club, 463, 479
- Coalition, 51
- Coalitional sovereignty, 35–36
- Coalition formation games, 39–40
- Coevolution, 725
- Cognitive agents, 895
- Collective/social choice problem, 529
- Collusive bidding, 346–347
- Combinators, 705
- Common intermediate language (CIL), 865
- Common prior, 129–31
 - and consistent beliefs, 119
- Common values, 345, 529
- (Pure) common values, 336
- Communicating finite-state machines (CFSM), 689
- Communication
 - direct, 795
 - equilibria, 181
 - indirect, 795
- Communication equilibrium and cheap talk, 114–116
- Communication network, 231
- Competitive equilibrium, 53
- Complete information, 185
- Completely revealing strategy, 158
- Complex adaptive system (CAS), 725
- Complex dynamics, 769
- Complexity, 913, 914, 923
 - and unanimity game, 647–651
- Complex system, 725
- Compositionality, 705, 706
- Computational algebra systems (CAS), 865, 874
- Computational intelligence (CI) principles, 793
- Computational mathematics systems (CMS), 865
- Computer generated imagery (CGI), 891
 - agent-based modelling in, 895–897
- Computer graphics technology
 - advances in, 891
 - agent-based modelling in, 892–895
- Computer languages, 865
 - design languages, 874–887
 - requirements, agent-based modeling, 871–874
 - types of, 869–871
- Computer programming language, 865
- Condition D1, 251
- Condorcet jury theorem, 529, 535
- Condorcet-Kemeny problem, 551–555
- Condorcet paradox, 529
- Condorcet's procedure, 550–551
- Condorcet winner, 543
- Configuration reachability (CR) problem, 694, 695
- Consistency assumption, 136
- Consistent beliefs, 129–31
- Consumer sovereignty (CS), 436
- Continuation game, 639
- Continuous time approach, 180
- Continuous-time games, 245–247
- Continuous two-sided matching model, 371
- Convergence, 488, 489
 - in classical learning models, 490
- Conway's game of life, 662

- Cooperative cost game
 - cooperative solutions, 427–429
 - cost sharing, 425–432
 - cost sharing rules, 432
 - definition, 425
 - dummy player, 425
 - egalitarian solution, 429–431
 - incentives in, 426–427, 442–444
 - nucleolus, 431–432
 - Cooperative games
 - definition, 49
 - extensive form, 50
 - representations of games, 50–51
 - theory, 50
 - Copeland's procedure, 555
 - Copycat strategy, 708–712
 - Core, 49, 52, 421, 463
 - Correlated equilibrium, 119, 229, 240–241
 - Bayesian equilibrium and, 134–136
 - in Bayesian game, 114
 - canonical representation, 109–110
 - and cheap talk, 114–116
 - and communication, 111–114
 - complexity, 110–111
 - definition, 109
 - duality and existence, 110
 - foundations, 111
 - geometric properties, 110
 - properties, 109
 - strategic form, 114
 - Correlation, 767
 - Correlation function, 780
 - Cost function, 421, 422
 - Cost of voting, 529
 - Cost sharing
 - additive, 441–442
 - Aumann-Shapley rule, 451
 - average rule, 440
 - Bayesian P^1 -demand games, 449
 - budget balance, 459
 - cardinal utility, 438
 - constant returns, 442
 - constrained stochastic, 457
 - continuity, 452
 - continuous homogeneous output model, 439–441
 - continuous models, 439–454
 - cooperative, 425–432
 - cooperative cost games, 425–432
 - definition, 421
 - egalitarianism, 457
 - Friedman-Moulin rule, 451
 - local symmetry, 457
 - marginal pricing rule, 441
 - monotone hazard condition, 439
 - Moulin-Shenker rule, 451
 - non-cooperative cost games, 432–439
 - nonlinear rules, 459
 - ordinality, 450
 - path-generated cost sharing rules, 450
 - positivity, 441
 - price of anarchy, 448
 - problem, 421, 423–425
 - production sharing problem, 422
 - proportional and serial equilibria, 446
 - rationing, 441–442
 - rule, 421, 423–424
 - rules induced by solutions, 432
 - scale invariance, 450
 - self-consistency, 452
 - serial rule, 440–441
 - Shapley-Shubik rule, 440, 452
 - stochastic models, 454–456
 - strategy-proofness, 449
 - uniqueness of nash equilibria, 446
 - upper bound, 452
 - Cost sharing problem, 422
 - Cournot best response dynamics, 486
 - Cournot duopoly, 33–34
 - Critical phenomena, 767, 778
 - Critical slowing down, 780
 - Cross-monotonicity, 437
 - Cryptographic model, 157
 - Customs and smugglers, 277
 - Cycle monotonicity, 326
 - Cyclic mapping, 913
- D**
- Deadlocking, 920
 - Decentralized control, 725
 - Decision-maker strategy, 168
 - Declarative language, 865
 - Dedicated agent-based modeling languages, 881
 - Demand games, 421, 434
 - mechanisms, 435–436
 - revelation, 435
 - in simple case, 445–446
 - straightforward, 437
 - Demand revelation game, 421
 - Design pattern, 866
 - Deterministic evolutionary dynamic, 573
 - Deterministic learning, 486
 - Deterministic macroscopic models, 853
 - Deterrence, 269
 - Differential games
 - antagonist games/zero-sum games, 306
 - definition, 305
 - dynamics, 306
 - features, 307
 - feedback form, 306
 - impulsive games, 311
 - with incomplete information, 313–314
 - Isaacs condition, 308
 - non antagonist, 311–312
 - objectives, 306
 - optimal value, 306

- payoff, 306
- qualitative and quantitative, 307
- qualitative target games, 307–308
- quantitative target games, 308–309
- rules of the games, 306
- stochastic, 312–313
- stochastic differential games, 312–313
- worst case design, 310–311
- zero sum, 309–311
- Diffusion equation, 774
- Digital organism, 725
- Digital organisms, 734
- Dimensionality
 - high, 772
 - low, 771–772
- Directed percolation, 785
- Direct interaction models, 675
- Discrete-event simulation, 685
- Discrete two-sided matching model, 371
- Discretization, 773
- Discriminatory stable sets, 13
- Dissemination of techniques, 923–924
- Dividend changes, 829
- Divinity, 251
- Dodgson's procedure, 556–557
- Domain's dimensionality, 323
- Domain-specific languages (DSLs), 664, 866, 875
- Domany-Kinzel model, 785
- Dominance relations, 612
- Dominant strategies, 317
- Domination, 9
- Downsian model of political competition, 529, 536
- Downward causation, 725, 733
- Dual game, 177–178
- Dummy, 428
- Duopoly market games, 33–34
- Dutch auction, 335, 336, 341
- Dutta's solution, 556
- Dynamic cellular-computing system, 791
- Dynamic method invocation, 866
- Dynamic network analysis, 726
- Dynamic programming principle, 309
- Dynamic/stochastic game, 65
- E**
- Echo, 726
- Economy, 463
- Efficiency, 499
- Efficient market believers (EMB), 835–836
- Egalitarianism, 529
- Egalitarian solution, 4298
- Elementary swarm, 798
 - See also* Swarm intelligence (SI)
- Embodied agent, 747
- Emergence, 726, 733
- Emergent features, 777
- Empirics of auctions, 343
- Encapsulation, 866
- Endogenous complexity, 644–646
- English auction, 335, 336, 338, 342, 346
- Entropy, 726
- Environment, 667
- Envy-freeness, 499
- Equilibrium, 61, 64–65, 229, 639
- Equilibrium binding agreements (EBA), 41
- Equilibrium concept, 349
- Equilibrium outcome, 251
- Equilibrium path, 639
- Equilibrium statistical mechanics, 775, 776
- Equitability, 499
- Equivalence principle, 53
- Equivalence theorem, 341, 343
- Essential superadditivity, 467
- Evolution (artificial), 726
- Evolutionarily stable strategy (ESS), 573
- Evolutionary algorithm, 726, 735
- Evolutionary computing, 726, 735
- Evolutionary game theory, 726
 - applications, 602–603
 - Bernoulli arrivals and mutations, 595–597
 - chaotic dynamics, 591–593
 - conservative properties, 586–590
 - definition, 574
 - deterministic dynamics, 580–592
 - global convergence, 586–590
 - local interaction, 599–602
 - local stability, 585–586
 - of Nash equilibrium, 583–585
 - with nonconvergent dynamics, 590–591
 - normal form games, 575–576
 - Poisson arrivals and payoff noise, 597–599
 - population games, 578–579
 - revision protocols, 579–580
 - static notions, 577–578
 - stochastic dynamics, 593–597
 - survival of dominated strategies, 592–593
- Evolutionary programming (EP), 736
- Evolution strategies (ES), 726, 735
- Excess volatility, 842
- Exchange economy, 53
- Existence, 65–67
- Exogenous complexity, 646–647
- Expected payoff, 339, 340, 344
- Extended system, 767, 769
- Extensive form, 269
- External stability, 9
- Extracellular matrix (ECM), 851
- F**
- Fair division
 - cake-cutting problem, 500–508
 - definition, 499
- Farsightedly consistent network, 632

Farsighted stable set, 9, 28–37
 applications, 30
 characteristic function form games and coalitional
 sovereignty, 35–37
 coalition formation games, 39–41
 cooperative games, 35
 in cooperative games, 35
 discrete public goods, provision of, 32–33
 duopoly market games, 33–34
 enforceability relation, 30
 house barter games, 43–45
 largest consistent set and largest farsighted
 conservative stable set, 29–30
 marriage games and roommate games, 42–43
 network formation games, 37–39
 n -person prisoner's dilemma, 32
 prisoner's dilemma, 31–32
 Feasible assignment, 371
 Feasible network, 614
 Feasible set network, 611, 617
 Fiber tracts, 851, 857
 Fictitious play, 487
 Field, 866
 Finite dynamical system, 683, 689
 agent-based simulation as, 692–693
 category of sequential dynamical systems, 699–700
 definition, 694
 mathematical results on, 696–698
 sequential update systems, 698
 as theoretical and computational tools, 693–694
 Finite field, 691
 Finite state machine (FSM), 726, 900
 First-generation models of computation, 706
 First-price auction, 335, 336, 338–341, 343
 First-price auction equilibrium, 340
 First-price sealed-bid auction, 338–340
 Flow payoff, 185
 Folk theorems, 243–244
 F-optimal stable matching, 372
 Formal model, 318–320
 Formation games, 626
 Friedman-Moulin rule, 451
 Frustration, 781
 Functional language, 866
 Function pointers, 866

G

Gale-Shapley algorithm, 376–377
 Game genres, 898–899
 Game horizon, 61
 Game of life, 662, 726, 791
 Games with incomplete information, *see* Bayesian games
 Game theory, 3–6, 49, 421, 639
 definition, 639
 dynamic games, 640
 General voting games, 14
 Generative social science, 726
 Generics, 866

Genetic algorithms (GA) model, 726, 735
 Genetic programming (GP), 726, 735
 Genotype, 726
 Gibbard-Satterthwaite theorem, 515–516, 529, 535–536
 Glioma cell migration
 cellular mechanisms of, 857
 fiber tracts on, 857
 Global climate change
 issues, 73–75
 models, 75–76
 results, 76
 Global pareto optimum (GPO), 76–77
 Glycolytic phenotype, 859
 Goto statement, 866
 Graph, lattice, tree, 767
 Grid, 913
 Grid computing, 923
 Groovy language, 885

H

Hamilton Jacobi Isaacs equation, 309, 310
 Handicap principle, 251
 Haptotaxis, 851, 857
 Harsanyi game revisited, 132–134
 Harsanyi's model, 121–122
 and hierarchies of beliefs, 122–124
 Headless, 866
 Heavy trading volume, 841–842
 Hedonic games, 40–41
 Heterogeneous directed network, 616
 Heterogeneous markets, 654–656
 Heterogeneous networks, 609
 Heuristic apparition of bimartingale, 174
 Hidden stochastic games, 180
 Higher-order programming, 866
 Homogeneous directed network, 616
 Homogeneous markets, 652–654
 Homogeneous networks, 609
 Hopfield networks, 688–689
 Hotelling model, 536
 House Barter games, 25–28, 43–45
 House exchange, 407
 Hybrid two-sided matching model, 372
 Hypercycles, 726, 733
 Hyperplastic phenotype, 859

I

Illegal production, 272–273
 ImmSim, 687
 Imperative language, 866
 Imperfect monitoring, 185, 241–243
 Implementable social choice rule, 349
 Implementation theory
 credibility, 362
 definition, 349
 dynamic implementation, 364
 environment, 351

- equilibrium, 352–353
- game-theoretical concerns, 361
- history, 350–351
- mechanisms, 352
- monotonicity, 357
- multiple, 363
- renegotiation, 362
- revelation principle, 354
- reverted preferences, 362
- social objectives, 352
- sociological factors/bounded rationality, 363–364
- Impossibility theorems, 511, 515
- Impulsive differential games, 311
- Impulsive games, 311
- Imputation, 9
- Incomplete information game, 185, 251
- Indirect domination, 9
- Indirect interaction models, 676–677
- Individual-based/interaction-based simulation, 684
- Individual-based modeling (IBM), 661, 726
- Inference approach, 344
- Information, 305
- Information economics, 287
- Information flow, 705
- Information processes, 732
- Inheritance, 866
- Inspection games, 4, 269
- Inspector leadership, 269
- Intelligence, 791
- Intelligent agent, 891
- Interaction, 667, 705
 - computation as, 706
- Interaction-based computing in physics, 768
- Interdependent valuations, 330
- Interdependent values, 337, 342
- Internal stability, 10
- Inter-node communication, 919
- Intuitive criterion, 251
- Isaacs condition, 308, 312
- Isaacs equations, 308
- Isaacs verification theorem, 309
- Ising model, 779
- Ising-type model, 913

J

- Jackson-van den Nouweland enforceability, 38
- Jackson-van den Nouweland rules, 620
- Jackson-Wolinsky enforceability, 38
- Jackson-Wolinsky network formation games, 633
- Jackson-Wolinsky rules, 611, 619
- Java, 866, 885
- Jointly controlled lotteries, 174

K

- Kalai-Lehrer learning, 491
- Kidney exchange, 413
- Kim and Markowitz portfolio insurers model, 827–828
- Ky Fan-Sion minimax theorems, 218–219

L

- Langton's ant, 726
- Langton's loop, 726, 731
- Large scale agent based models, 917
- Large scale simulations
 - benefits, 921
 - definition, 914
 - pitfalls, 921–922
- Largest consistent set (LCS), 30
- Lattice Boltzmann equation, 784
- Lattice-gas cellular automata (LGCA), 784, 855, 857
- Lattice property, 372
- Learning
 - definition, 485
 - sophisticated, 492–493
 - stochastic, 495–496
- Learning classifier system (LCS), 726, 735
- Lenz-Ising model, 780
- Liberalism, 511, 514, 515, 520
- Lindenmeyer system (L-system), 726, 732
- Linear bounded automaton (LBA), 695
- Linear logic, 705
- Linear systems, 772
- LLS model, 830–835
- Load balancing, 919
- Logic and geometry of interaction, 708
- Logic programming language, 866
- Long-lived player, 185
- Lux and Marchesi model, 829–830
- Lux model, 829–830
- Lyapunov exponent, 783

M

- Machine learning, 726
- Macro language, 866
- Macrovariables, 778
- Majority relation, strict, 543
- Majority voting, 529
- Manipulable mechanism, 372
- Many particle systems, 770
- Marginal cost pricing, 441
- Market anomalies, 825
- Market clearance, 836–837
- Market design
- Market games, 463
 - equivalence, 469–472
- Markov approximation, 774–775
- Markov chain games, 178
- Markov chains, 774
- Markov-perfect equilibrium, 77–78
- Marriage games, 20–25, 42–43
- Massively multi-agent systems (MMAS), 664, 914
- Matching μ , 372
- Matching mechanism, 372
- Matching theory
 - applications, 411–417
 - one-sided matching, 406–411
 - two-sided matching, 401–406

Mathematica, 732, 866, 877, 878
 Mathematical framework, 683
 MATrix LABoratory (MATLAB), 866, 877, 879–881
 Maximality, 499
 Maximin preferences, 372
 Maynard Smith replicator dynamic, 582
 McCoy's equilibrium strategy, 282
 Mean-field approach, 767, 771, 775
 Medical residency matching, 412
 Meme, 726
 Message passing (MP), 914
 Message passing interface (MPI), 914
 Metabolism, 857–858
 Method, 866
 Methodological individualism, 866
 Microlevel simulation, 819
 Microscopic simulation, 826
 Microsoft intermediate language (MSIL), 865
 Mind-body problem, 726, 734
 Minimax theorem
 and statistical estimation, 220–221
 zero-sum two person games, 217–218
 Mixed strategy, 269
 Mobile agents, 866
 Mobilization and group-based notion of welfare, 538
 Model analysis, 923
 Modeling physical systems, 771
 Module, 867
 Molecular dynamics, 772
 Money left on the table, 345
 Monitoring structure, 139
 Monte Carlo method, 767
 Monte Carlo simulations, 776
 Moral hazard (hidden action), 287
 Morphological computation, 747
 Moulin-Shenker rule, 451
 Multi-agent systems (MAS), 661
 Multi-dimensional domains, 325–326
 Multiplayer games, 237–240
 Multiple instruction multiple data (MIMD), 914
 Myopic optimum, 185

N

Nash equilibrium, 107, 139, 160, 185, 251, 269, 433
 Nash reversion, 185
 Natural asynchrony, 791
 Neologism-proof equilibrium, 251
 NetLogo, 867, 881
 Network formation games, 37–38
 Networks, 778
 abstract games, 624–626
 Bala-Goyal rules, 611
 definition, 610
 directed, 616
 dominance relations, 612
 farsightedly consistent, 632
 feasible, 614
 feasible set, 611, 617

 formation, 619
 formation games, 626
 games of network formation, 610
 heterogeneous directed, 616
 homogeneous directed, 616
 Jackson-van den Nouweland rules, 620
 Jackson-Wolinsky rules, 611, 619
 paths and circuits, 618
 payoff functions, 619
 in physics, 778
 players' preferences, 611, 618–619
 rules of network formation, 611
 rules supernetworks, 621–622–624
 strongly stable homogeneous, 630–632
 supernetworks, 622
 types of, 614–615
 without descendants, 626–630
 Neumann-Morgenstern stable sets, 28
 Newton's laws, 770–771
 No excess volatility, 838–841
 Noise, 770
 Non-cooperative cost games, 432–437
 Non-cooperative implementation, 54–55
 Non-emptiness, 52–53
 Nonlinear element, 768
 Nonlinear system, 767
 Non-player character (NPC), 891
 Non-procedural language, 867
 Non-transferable utility (NTU) game, 12, 36, 51
 Non zero-sum games
 biconvexity and bimartingales, 175–176
 equilibrium payoffs, 171–175
 existence of equilibria, 171–173
 No-regret strategies, 167
 Normal form game, 573
 Normalized discounted value, 185
 Normal modes, 772
N-Person prisoner's dilemma, 31
 Nuclear material diversion, 274–277
 Numerical physics, 769

O

Object, 867
 Objective-C, 867
 Object-oriented languages, 867, 872
 Observer, 867
 ODD protocol, 867, 874
 Off-the-equilibrium-path play, 656
 Ontogenesis, 747
 Optimization algorithms, 791
 Ordinary differential equations (ODEs), 685
 Out-of-equilibrium phenomena, 785

P

Pairwise majority voting rule, 532
 Parallel computer architecture, 914
 Parallel computing, 918–919

Pareto-optimal matching, 372
 Partial differential equations (PDEs), 685, 773
 Partially revealing strategy, 158
 Particle swarm optimization (PSO), 726, 738, 791, 799–801
 Partition function, 776
 Passenger ticket control, 271
 Patch, 867
 Payoffs, 61
 feasible, 70–72
 functions, 619
 individually rational, 72–73
 uncertainty, 493–494
 vector, 12, 463, 466
 Per capita boundedness, 463, 472
 Percolation, 767
 Perfect monitoring, 185
 Performance moderator functions (PMFs), 906
 Perron's theorem on positive matrices, 213–214
 Phase transitions, 777
 Phenotype, 726
 Pheromone, 792
 Phylogenesis, 747
 Platform for agent-based simulation, 667
 Players, 61
 Players' preferences, 611, 618–619
 Plurality rule, 549–550
 with runoff, 549–550
 Plurality voting rule, 533
 Political competition, 536
 Political ignorance, 535
 Polymorphism, 867
 Pooling equilibrium, 251
 Population game, 573
 Power set, 421
 Pre-auction knock-out, 346
 Precision of information, 342
 Prediction rule, 485
 Preference, preference aggregation, 543
 Pregame, 463
 Premarkets, 471
 Price taking equilibrium, 464
 Principal agent, 287
 Principal-agent models
 adverse selection model, 294–299
 base game, 290–291
 contracts and social preferences, 299–300
 definition, 287–288
 empirical studies of, 299
 moral hazard model, 291–293
 repeated moral hazard, 293
 Prisoner's dilemma, 31
 Private values, 336
 Probabilistic cellular automata, 785
 Probability
 chaos and, 773
 distribution, 767, 773
 Procedural language, 867
 Production market games, 14–15

Profile, 543
 Programming language, 867
 Projection operation, 772
 Proportionality, 499
 Pure strategy (simply strategy), 107

Q

Qualitative target games, 307–308
 Quantitative target games, 308–309
 Quantum teleportation, 721
 Quasi-linear utilities and VCG mechanism, 321–323
 Quota of agent, 372

R

Random access memory (RAM), 918
 Random walk, 773–774
 Reaction-diffusion system, 726, 732
 Receiver, 251
 Record, 867
 Recursively generated object, 727, 732
 Recursive porous agent simulation toolkit (Repast), 867, 881
 Reduction of complexity, 770
 Reflection, 867
 ReLogo, 867
 Rematching proof equilibrium, 373
 Repast, 678
 Repast symphony, 867, 885
 flowcharts, 881
 Repeated game, 139, 185, 485, 639
 almost-perfect monitoring, 145–146, 153
 almost-public monitoring, 153
 canonical structure, 146
 communication equilibria, 147
 correlated equilibria, 147
 definition, 139
 deterministic signals, 146
 entropy characterizations, 154–155
 equilibrium notions, 141
 finitely, 143
 folk theorem, 143
 folk theorems, public equilibria, 152
 impatient players, 142–143
 induced plays, 140
 infinitely patient players, 142
 nash equilibria, 146
 nash reversion, 143
 necessary conditions, equilibrium payoffs, 141
 with non-observable actions, 144–145
 with observable actions, 140–144
 observable payoffs, 146
 perfect monitoring, 145
 preferences, 140
 prisoner's dilemma, 139, 143
 private strategies, 152
 public and almost-public signals, 146
 public communication, 152

- Repeated game (*cont.*)
 - public equilibria, 150–151
 - punishment levels, 154
 - recursive structure, 143
 - sequential equilibria, 146
 - with unique nash payoff, 144
 - Repeated game model
 - dynamic interaction, 62
 - dynamics, 73
 - folk theorem, 70
 - histories and strategies, 63–64
 - Repeated game with incomplete information
 - Aumann and Maschler model (*see* Aumann and Maschler model)
 - definition, 158
 - equilibria of general-sum games, 160
 - non-observable actions, 175–177
 - payoffs, 159–160
 - perturbations of, 181
 - strategies, 159
 - vector payoff and approachability, 164–169
 - zero-sum game values, 160
 - Replicator dynamic, 573
 - Representative agent, 825
 - Reputation bound, 185
 - Reputation effects
 - complete information, 190
 - definition, 185
 - H-action type, 189
 - Nash equilibrium, 191
 - product-choice game, 187
 - pure-action Stackelberg payoff, 192
 - stage game, 190
 - temporary reputation effects, 194
 - two long-lived players, 197–198
 - Responsive preference, 373
 - Return autocorrelation, 845
 - Revelation game, 373
 - Revelation mechanism, 373
 - Revenue maximizing auction, 342
 - Revision protocol, 573
 - Road network representation, 686
 - Roommate games, 42–43
 - r(f)-separable preference, 372
 - Runtime type identification (RTTI), 867
- S**
- Saddle point, 269
 - Scale-free networks, 778
 - Schelling's segregation model, 826–827
 - Second-order phase transition, 780
 - Second-price auction, 335–338
 - Selection bias, 343
 - Self-organization, 727
 - Self-organized criticality, 778
 - Self-replication, 727
 - Selfridge-Conway procedure, 502
 - Self-signaling set, 251
 - Semipermeable barriers, 308
 - Sender, 251
 - Sensibility to variation, 772
 - Separating equilibrium, 251
 - Sequential dynamical system (SDS), 690, 694
 - Sequential equilibrium, 107, 139
 - Shapley-Shubik rule, 440, 452
 - Shapley value, 49, 56, 427–428, 464
 - Sharing common pool resources, 280
 - Short-lived player, 185
 - Signaling game, 4, 251
 - applications, 262–265
 - basic model, 254–256
 - in biology, 263–264
 - cheap talk, 257–261
 - communication about intentions, 261–262
 - definition, 252
 - economic applications, 262
 - equilibrium, 253–254
 - model, 253
 - multiple equilibria and selection, 256–257
 - political science, 264–265
 - verifiable information, 261
 - Simple action type, 185
 - SIMULA (SIMULATION LAnguage), 872
 - Simulated annealing, 776
 - Sincere strategy, 373
 - Single-crossing condition, 251
 - Single-dimensional domains, 323–325
 - Single instruction multiple data (SIMD), 914
 - Single peakedness, 511, 516, 517, 520
 - Single-peaked preferences, 546
 - Single-tape TM model, 695
 - Singleton basins, attraction, 632–634
 - Situated agent, 747
 - Slater problem, 551–555
 - Slime trail, 855
 - Slime trail motion, 851
 - Small group effectiveness (SGE), 464, 473
 - Small world effect, 778
 - SOAR architecture, 904
 - Social ability, 661
 - Social agent-based modeling, 727
 - Social choice correspondence, 543
 - Social choice function, 317, 512, 521, 522, 543
 - definition, 536
 - manipulability, 529
 - Social choice rule, 349
 - Social choice theory
 - broader horizons, 523–524
 - cyclical patterns and Arrow's impossibility theorem, 512–514
 - definition, 512
 - escape routes, 516–521
 - Gibbard-Satterthwaite theorem, 515–516
 - paretian liberal, 514–515
 - voting rules, 521–523
 - Social network analysis (SNA), 727

- Social phenomena simulation
 - beliefs, desires, and intentions (BDI) model, 822
 - definition, 819
 - dynamic microsimulation, 822
 - engineering of systems, 821
 - polymaking, 821
 - supporting planning, 821
 - training, 821–822
 - types, 820
- Solution concept, 49, 52
- Somatic evolution, 851, 859
- Spence-Mirrlees condition, 251
- Stable matching mechanism, 373
- Stable outcome, 373
- Stable sets, 11
 - abstract games, 11
 - applications, 13
 - assignment games, 17–20
 - characteristic function form games, 12
 - and core, 11, 12
 - farsighted, 28–37
 - farsighted conservative, 29
 - game with no stable set, 12
 - general voting games, 14
 - house barter games, 25–28
 - marriage games, 20–25
 - production market games, 14–17
 - symmetric voting games, 13–14
- Stackelberg action, 185
- Stackelberg type, 185
- Stage game, 185
- Stage payoffs, 229
- Stand-alone cost game, 425
- Standard signaling game, 252
- State independent signalling, 176
- State of nature, 119
- State of system, 767
- State of world, 119
- State space, 768
- Static game, 308, 309
 - asymmetric information version, 103–104
 - battle of sexes, 84
 - Bayesian games, 102–103
 - chicken models, 84
 - correlated equilibrium, 94–98
 - definition, 83
 - equilibrium in beliefs, 92–94
 - evolutionary stable strategies, 104–105
 - existence, 87–88
 - matching pennies, 84, 86
 - mixed strategies, 88–91
 - nash equilibrium, 85
 - prisoner's dilemma, 84–86
 - rationality and correlated equilibrium, 101–102
 - war of attrition, 86, 91–92
- Statistics, 773
- Stigmergy, 727, 737, 792, 795
- Stochastically stable state, 573
- Stochastic differential equations, 774
- Stochastic differential games, 312–313
- Stochastic finite dynamical systems, 691–692
- Stochastic game, 159
 - algorithms, 244–245
 - capital accumulation, 230–231
 - communication network, 231
 - continuous-time games, 245–247
 - correlated equilibrium, 240–241
 - definition, 229
 - equilibria, 231–233
 - evaluations, 231–233
 - folk theorems, 243–244
 - imperfect monitoring, 241–243
 - multiplayer games, 237–240
 - queues, 231
 - strategies, 231–233
 - taxation, 231
 - zero-sum games, 233–237
- Stochastic optimization, 776–777
- Stochastic SDS (SSDS), 696
- Strategic complexity, 639
- Strategic demand games, 432
- Strategic (normal) form game, 10, 107
- Strategic game, 421
- Strategic situations, 50
- Strategic voting, 536
- Strategies, 61, 229, 305, 639
- Strategy-proofness, 436, 512, 520
- Strongly stable homogeneous, 630–632
- Strongly substitutable preferences, 373
- Structured language, 867
- Subgame, 185
- Subgame perfect equilibrium, 185
- Substitutable preferences, 373
- Sugarscape, 727
- Super-agents, 916
- Supernetworks, 622
- Swarm intelligence (SI), 663, 727, 736, 792, 867, 883–885
 - advantages, 793
 - ant colony optimization (ACO), 801–802
 - applications, 793
 - artificial life systems, 797–798
 - asynchronous, 807–811
 - biological systems, 794–795
 - cellular-computing methods, 803–804
 - characteristics, 792, 811–812
 - definition, 792, 798–799
 - difficulties in, 794
 - dynamics, 812
 - features, 793
 - implicit assumption, asynchrony irrelevance, 806–807
 - importance, 793
 - intelligence, universal computation, 804
 - intelligent units, 813–814
 - intuitive notion, 793
 - investigations, 793
 - limitations, 803
 - local synchrony and self-synchronization, 809
 - modeling asynchrony, 807–809

Swarm intelligence (SI) (*cont.*)
 natural asynchrony, 809–810
 nonlinear differential equation methods, 802
 optimization, 800
 particle swarm optimization (PSO), 800–801
 randomness in, 805–806
 realization of asynchronous, 810–811
 relations to standard-mathematics methods, 804–805
 robotic systems, 795–797
 standard-mathematics methods, 799, 804–805
 swarm optimization, 799–800
 term usage, 792
 types of asynchrony, 807
 unpredictability, 812–813
 Swarm optimization, 792
 Swarm robotics, 792
 Symbolic processing, 727, 732
 Symmetric information, 180
 Symmetric voting games, 13
 Symmetry, 428
 Synchronization, 783

T

Tarski's theorem, 384
 Taxation, 231
 Templates, 867
 Theoretical models, 769
 Tierra, 727
 Time-averages, 770
 Time-series, 774
 Time-series data, 774
 Time-stepped simulation, 685
 T-map, 373
 Totally balanced game, 464
 Traffic and crowd agent models, 668
 Trajectory, 768
 Transferable utility (TU) game, 12, 36, 51
 TRANSIMS (Transportation, Analysis and Simulation System), 685, 686, 696
 Truncation, 373
 Tumor invasion
 cell-cell and cell-ECM adhesion, 856
 effects of directed cell motion, 855–856
 emergence of, 858–859
 game of invasion, 860
 genetic alterations, 853
 heterogeneous environments, 857
 mathematical modeling and analysis, 853
 metabolic changes influence, 859–860
 migration, 856
 models of, 855
 morphology, 855
 phase, 852
 processes, 852
 spatial structure of, 856
 Turing machine (TM), 694
 Turtle, 867
 Two-EMB-population model, 843

Two-sided matching model, 373
 with additively separable utility functions, 386–392
 definition, 373
 discrete, 378–385
 Gale-Shapley algorithm, 376–377
 historical account, 375–376
 hybrid one-to-one, 392–394
 incentives, 394–397
 Type (state of mind), 119

U

Umbrella folding algorithm, 213
 Unified modeling language (UML), 867, 875
 Universal belief space, 124–125
 Universal convergence, 491–492
 Universal Turing machine (UTM), 727, 734
 Unpredictability, 775
 Unpredictable system, 792
 Unstructured language, 868
 Update rule, 727
 Utility function, 107, 269

V

Value functions, 309
 VCG mechanisms, 317
 Vector computer/vector processor, 914
 Vehicle dynamics, 686
 Verifiable information, 261
 game, 252
 Verification, 269
 Viability theory, 310
 ViCrowd system, 895
 Virtual character, 891
 Virtual fidelity, 895
 Virtual machine, 868
 Voluntary participation (VP), 436
 Von Neumann architecture, 792
 Von Neumann machine, 661
 Von Neumann-Morgenstern stable set, 9, 10
 definition, 10
 von Neumann-Morgenstern utility function, 107
 Von Neumann's self-replication problem, 729
 Voting
 Borda Count, 533
 common value setting with strategic agents, 539
 maximin rule, 534
 methods, 5, 512, 524
 plurality, 533
 supramajority, 533
 unanimity, 533
 Voting cycle, 532
 Voting procedures
 acyclic relation, 548
 approval voting procedure, 557
 Arrow, Kenneth J., 545
 Borda, Jean-Charles de, 544

Borda's procedure, 550
 Bucklin's procedure, 558
 Caritat, Marie Jean Antoine Nicolas de, 544
 complete preorder, 548
 complexity classes, 558–564
 Condorcet-Kemeny problem, 551–555
 Condorcet-Kemeny winner, 552
 Condorcet, marquis de, 544
 Condorcet's procedure, 550–551
 Copeland's procedure, 555
 definition, 543
 Dodgson's procedure, 556–557
 Dutta's solution, 556
 Lhuillier, Swiss Simon, 545
 majority tournament, 548, 551
 maximin procedure, 554
 median orders, 551–555
 median relation, 552
 minimax procedure, 554
 plurality rule, 549–550
 plurality rule with runoff, 549–550
 preferential, 549–550
 preorder, 548
 prudent order, 554
 ranked pairs procedure, 554
 Slater problem, 551–555
 symmetric difference distance, 552
 tournament solutions, 555
 variants, 558
 Young's procedure, 557
 Voting system, 687
 Voting theory, 543

W

Walrasian equilibrium, 651
 Weak monotonicity, 326
 Welfare economics, 533–534
 Winner's curse, 344–346

Y

Young's procedure, 557

Z

Zero-sum games, 160, 233–237, 269
 correlated initial information, 170
 differential games, 309–311
 limit values, 170
 maxmin and minmax of repeated, 169–170
 operator approach, 180
 uncertain duration, 180
 Zero-sum two person games
 adjacent vertices, 202
 algorithm application, 204–205
 approximate fixed points, 203–204
 behavior strategies, 215–216
 Borel's Poker model, 221–223
 fictitious play, 209–210
 game of Hex, 202–203
 games with perfect information, 200
 games with perfect recall, 214–215
 geometric consequences, 218
 historical remarks, 207
 infinite games applications, 219–220
 Ky Fan-Sion minimax theorems, 218–219
 minimax theorem, 217–218
 minimax theorem and statistical estimation, 220–221
 mixed strategy and minimax theorem, 206–207
 nim game, 200–202
 optimal strategies via linear programming, 207–208
 payoffs, 223–225
 Perron's theorem on positive matrices, 213–214
 saddle point, 205–206
 search games, 210–213
 sequence form, 216–217
 S-games, 217–218
 simplex algorithm, 208–209
 umbrella folding algorithm, 213